arXiv:1612.08699v1 [stat.ME] 27 Dec 2016

Comparative Causal Mediation: Relaxing the Assumption of No Mediator-Outcome Confounding

Kirk Bansak*

December 2016

For comments and questions,
please e-mail: kbansak@stanford.edu

Abstract

Experiments often include multiple treatments, with the primary goal to compare the causal effects of multiple treatments. While comparing the magnitudes of the average treatment effects (ATEs) is straightforward, there exist few methods to systematically compare the causal anatomies of each treatment—that is, the collection of causal mechanisms underlying each treatment's total effect—in order to understand the sources of their relative magnitudes. This study introduces a framework for comparing the causal anatomies of multiple treatments through the use of causal mediation analysis. The study proposes a set of comparative causal mediation estimands that compare the mediation effects of different treatments via a common mediator. It derives the properties of a set of estimators, which are shown to be consistent (or conservative) without making the assumption of no unobserved confounding of the mediator-outcome relationship, which is a strong and nonrefutable assumption that must be made for consistent estimation of individual causal mediation effects. The estimators are easy to understand and implement, thereby providing researchers with a simple, reliable, and systematic method of comparing, discovering, and testing the causal mechanism differences between multiple treatments. An original application is presented to illustrate the method.

Keywords: causal mediation, causal inference, experimental design

*Department of Political Science, Stanford University, 616 Serra Street, Stanford, CA 94305-6044.

I thank Avidit Acharya, Justin Grimmer, Jens Hainmueller, Andy Hall, Kosuke Imai, Hye-Sung Kim, Ken Scheve, Mike Tomz, and Teppei Yamamoto for helpful comments.

I. Introduction

Causal mediation analysis aims to open the "black box of causality," offering the opportunity to explore how and why certain treatment effects occur in addition to simply detecting the existence of those effects. Estimation of causal mediation effects, which are effects transmitted via intermediary variables called mediators, has traditionally been implemented using the parametric structural equations model (SEM) framework (Baron and Kenny, 1986). Recent years have seen important advances in the formalization, generalization, and estimation of causal mediation effects within the potential outcomes framework (Imai et al., 2010a, 2011a,b) and nonparametric SEM framework (Pearl, 2001). Increasingly, causal mediation analysis is also used in experimental research. In the most commonly used causal mediation experimental design, sometimes called "the single-experiment design," the treatment variable is randomized and the mediator(s) observed.

Another trend in experimental research is the design of more complex experiments featuring multiple treatment arms. As knowledge and empirical results have accumulated in various academic sub-fields and in specific program evaluation contexts, experimental research questions have evolved in ways that require evaluating multiple related treatments. Instead of simply testing the effects of single treatments, often of primary interest are the empirical and theoretical differences between the effects of multiple treatments. Indeed, scientific and social scientific theories usually entail causal mechanisms, and hypotheses on why one treatment should have a larger effect than another tend to be based on the mechanism(s) through which each treatment propagates its effect. In such research contexts, richer insights can be gained from comparing the different treatments' causal anatomies—that is, the ensemble of causal mechanisms that endow each treatment with its effect.

This paper presents a new framework for comparing the causal anatomies of multiple treatments through the use of causal mediation analysis. It proposes a novel set of comparative causal mediation (CCM) estimands that compare the mediation effects of different treatments via a common mediator. In addition, the value of the method is enhanced by the fact that, as this paper shows, these CCM estimands can be estimated under fewer threats to internal validity than individual causal mediation effects. Specifically, consistent estimation of individual causal mediation effects requires the strong and nonrefutable assumption of no unobserved confounding of the mediator-outcome relationship. In contrast, this paper derives the properties of a set of estimators for the CCM estimands and shows these estimators to be consistent (or conservative) without making the assumption of no unobserved confounding of the mediator-

outcome relationship. The paper also derives the finite-sample properties of these estimators, showing how adjustments can be made to improve the central tendencies of these estimators in small samples.

The estimators are easy to understand and implement, thereby providing researchers with a simple, reliable, and systematic method of comparing, discovering, and testing the causal mechanism differences between multiple treatments. An original application is presented to illustrate the method. Software is also offered to implement the method.

II. Motivation for Comparing Causal Mediation Effects

As research agendas in both academic and policy contexts advance, experimental research often involves evaluating distinct but conceptually or administratively related treatments. Comparing the magnitudes of the average treatment effects (ATEs) of multiple related treatments is a valuable first step, but simple ATE comparisons provide little insight into the specific processes by which one treatment achieves a larger effect over another—and, alternatively, whether equally sized treatments are truly interchangeable or instead work via different channels. In research contexts where multiple related treatments are under study, richer insights can be gained from comparing the related treatments' causal anatomies. Such investigations should be useful for both theoretical discoveries and policy/program evaluation purposes.

A. Testing Theories

While experimental tests of social scientific theories often begin and end with estimation of ATEs, those theories almost always also involve causal mechanisms or channels that remain untested. And in research contexts involving multiple related treatments, theories on why one treatment should have a larger effect than another are based on the mechanism(s) through which each treatment propagates its effect. One example exists in the well-developed literature on party cues in American politics, which includes a wealth of experimental studies that investigate party cue effects on voter attitudes and behavior.[1] As this literature has highlighted, not all party cues are made equal; instead, there are various types of party cues, with empirical evidence suggesting that some work more forcefully than others. For instance, there is some experimental evidence that out-party cues may, in fact, be more influential than in-party cues (Aaroe, 2012; Arceneaux and Kolodny, 2009; Slothuus and de Vreese, 2010; Goren et al., 2009;

[1]Party cues are public signals from political parties that associate a party with particular candidates or policy positions, thereby affecting the attractiveness of those candidates or positions for voters who have partisan orientations.

Nicholson, 2012). Yet while there is evidence on the relative magnitudes of these two types of cues, the reasons or precise mechanisms by which out-party cue effects exceed those of in-party cues have been theorized but not tested in a disciplined manner. Doing so would involve a rigorous method for comparing the mechanisms underlying each set of party cues.

As another example, consider the recent accumulation of survey experimental evidence on audience costs in international relations (for a brief review, see Hyde, 2015). Audience costs refer to the electoral costs to politicians (i.e. punishment by voters) for breaking policy commitments, and the past ten years have seen a deluge of survey experiments providing evidence that voters do, indeed, tend to punish policymakers for reneging on foreign policy commitments. These many studies have differed greatly, however, not only in terms of their foreign policy contexts (e.g. security scenarios, international economic scenarios, etc.) but also in terms of the specific nature of the foreign policy commitment (e.g. informal, legal, etc.). One may then wonder whether and why the nature of such a commitment might affect the strength of the audience cost effect—for instance, through what mechanism may a legalized foreign policy commitment gain audience cost strength over an informal commitment. This is a question that can be answered by comparing the causal anatomies of different audience cost treatments, which is explored in the application presented in section VII of this paper.

B. Policy and Program Evaluation

Comparing the causal anatomies of related treatments also offers great value in the policy and program evaluation context, where multiple related treatments are often investigated in individual studies. The basic goal in this context is to identify, often through experimental studies, which policy interventions are effective in achieving certain economic, social, political, or other outcomes. Ideally, the effectiveness of any preferred policy intervention should be generalizable across time and different localities. However, because of constraints on resources, as well as logistical and administrative realities, the execution of experimental studies is often restricted to short periods of time and small subsets of locations. This results in empirical findings with great internal validity in their localized context but potentially questionable external validity, limiting the researcher's confidence on whether the treatment found to be most effective in the study would actually be the most effective one more widely.

Ultimately, the goal of such studies is to enable policymakers (or other actors) to select treatments that will be effective beyond the initial study locality, and this requires evaluating how different treatments are likely to perform in different contexts. One important means of doing so is developing a comprehensive understanding of the mechanisms underlying different

treatments, as a treatment's performance in a particular context depends upon whether or not its primary mechanism(s) can properly operate in that context.

For instance, consider an experimental study on job training programs, aimed at finding employment for unemployed lower-income adults. Imagine that the study is implemented in a handful of towns and involves multiple training programs under consideration; for simplicity, assume two training programs (i.e. two treatments and a control condition of no training). A preliminary analysis of the results may reveal, for instance, that both programs have roughly equal-sized effects on employment. A superficial interpretation of these results would then be that the two programs are interchangeable and hence, upon launch of a follow-up large-scale policy intervention, the choice of which specific training program to roll out could be made arbitrarily or based on other criteria. However, given that the experimental results of the study were based on a single sample in a single snapshot in time, such a decision-making rule could be severely sub-optimal.

Instead, it would be useful to investigate the causal mechanism differences between the two job training programs. While their ATEs may have been roughly equal, it is possible that they achieved their positive effects on employment via different channels. One program may have achieved its primary effect by increasing the job search motivation of its participants, while the other may have achieved its primary effect by helping its participants to develop specific technical skills. If equipped with such knowledge, policymakers would be in a much better position to make optimal decisions on which job training program to introduce in different localities. For instance, certain local economies may have a low supply of low-skill jobs and a relatively higher supply of technical jobs, and thus the training program working primarily through the skill-development channel would be preferred in this locality. In other words, knowledge of the relative causal anatomies of the two training programs would enable more insightful comparative evaluation and ultimately more optimal policy intervention decisions.

In sum, much is to be gained from both a theoretical and practical policy standpoint from evaluating the causal anatomy differences between treatments. The method of comparative causal mediation presented in this paper provides a new, simple, and rigorous means of doing so.

III. Comparative Causal Mediation (CCM) Estimands

Let the average causal mediation effect (ACME) be the average effect of a treatment T on an outcome Y transmitted via a mediator M.[2] As will be shown, in spite of the strong assumptions required for the identification of any single ACME, a weaker set of assumptions—which, notably, does not contain the usual assumption of no unobserved confounding of the mediator-outcome relationship—allows for the identification of the following two comparative causal mediation (CCM) estimands of interest.

Definition 1: *Define the estimands of interest as follows:*

$$Estimand\,1: \ \frac{ACME_{T=t_2,M=m}}{ACME_{T=t_1,M=m}}$$

$$Estimand\,2: \ \frac{(\frac{ACME_{T=t_2,M=m}}{ATE_{T=t_2}})}{(\frac{ACME_{T=t_1,M=m}}{ATE_{T=t_1}})}$$

where t_1 and t_2 denote two mutually exclusive treatments, m denotes a common mediator, and all effects ($ACMEs$ and $ATEs$) are referenced against a control condition (i.e. a third condition $c \neq t_1 \neq t_2$).

For conceptual clarity, consider a three-level experimental design that involves a true control condition and two different treatments. The two treatments may be qualitatively different or one may be a scaled up version of the other. Furthermore, there is a single mediator of interest. It may be the case that multiple mediators have been measured in the experiment, but the estimands are applied within the context of a single mediator at a time. Hence, the estimands can be re-expressed more simply as:

$$Estimand\,1: \ \frac{ACME_2}{ACME_1}$$

and

$$Estimand\,2: \ \frac{(\frac{ACME_2}{ATE_2})}{(\frac{ACME_1}{ATE_1})}$$

where, for instance, $ACME_2$ refers to the average effect of treatment 2 (relative to the control condition) that is mediated through the mediator of interest.

The first estimand measures the extent to which treatment 2 has a greater causal mediation effect than treatment 1 in terms of the mediator of interest. This allows for testing whether one

[2]This is often also referred to as the indirect effect, with terminology differing across researchers and depending upon the statistical framework being employed.

treatment has a stronger absolute effect transmitted through the mediator of interest, relative to another treatment. For instance, in the hypothetical job training policy evaluation described above, it could measure whether one job training program had a stronger effect transmitted via job search motivation than the other program. In contrast, the second estimand measures the extent to which treatment 2 has a greater proportion of its total effect transmitted through the mediator of interest, relative to treatment 1. This allows for testing the extent to which the mediator is more important to the overall causal anatomy of one treatment versus another. Again, in the context of the job training policy evaluation, this could measure whether effect transmission via job search motivation comprised a larger proportion of the overall effect of one program versus the other. Table 1 summarizes the general research questions related to each CCM estimand.

Table 1: General Research Questions Related to Each CCM Estimand

Estimand 1	Does ATE_2 exhibit stronger effect transmission via mediator M than ATE_1 does? Does the second treatment have a larger mediated effect in absolute terms? $H_0 : \frac{ACME_2}{ACME_1} = 1 \quad H_a : \frac{ACME_2}{ACME_1} > 1$
Estimand 2	Does effect transmission via mediator M make up a larger proportion of ATE_2 relative to ATE_1? Is M more important for ATE_2 than ATE_1? $H_0 : \frac{\left(\frac{ACME_2}{ATE_2}\right)}{\left(\frac{ACME_1}{ATE_1}\right)} = 1 \quad H_a : \frac{\left(\frac{ACME_2}{ATE_2}\right)}{\left(\frac{ACME_1}{ATE_1}\right)} > 1$

Which of the two estimands is of interest will depend upon the empirical and theoretical goals of a particular research project. When the researcher's main goal is to identify which treatment has the strongest absolute effect transmitted via a specific causal channel, the first estimand is likely to be of primary interest. The case of evaluating different job training programs, as presented earlier, provides an example. From the standpoint of optimal policy implementation, the researcher may choose to focus on one specific causal channel, prioritizing transmission of the causal effect via that channel and discounting transmission via other channels. For instance, if the researcher knows that the training programs under consideration will, in the post-evaluation period, be rolled out in target areas where increasing job-search motivation is unlikely to be an effective method of increasing employment (e.g. in

local economies with a low supply of low-skill jobs), then it makes sense for the researcher to prioritize the skill-development causal channel. In other words, the researcher's goal should be to identify which job training program leads to the largest increase in employment specifically via the skill-development channel, regardless of the magnitude of the effect transmitted via the channel of job-search motivation and perhaps even regardless of the relative magnitudes of programs' overall ATEs. In that case, the researcher's goal would be achieved by investigating the first CCM estimand, which would measure how much larger one treatment's skill-development causal channel is than that of the alternative treatment(s).

If, instead, the researcher is interested in better understanding multiple treatments' relative causal anatomies more generally, then both the first and second CCM estimands should be of interest. Considering both estimands could be useful in particular for theoretically motivated researchers who are seeking to test theories involving multiple treatments. Such theories not only predict whether one treatment should be more effective than another but also often dictate (a) the specific causal mechanisms that should grow or shrink when switching from one treatment to another and (b) the specific causal mechanisms that should contribute a larger share of the overall ATE for one treatment versus another. Indeed, for the purposes of theory testing and exploration, the two CCM estimands could be considered in conjunction with the ATEs to form a full picture of the relative causal anatomies of different treatments. To illustrate, Table 2 provides a set of some of the theoretical implications that would follow from testing hypotheses about the CCM estimands in combination with the ATEs.

IV. Comparative Causal Mediation Estimation Framework

Estimators of the CCM estimands, as presented in this paper, will be formulated within the parametric structural equations model (SEM) framework. A number of researchers have critiqued the parametric SEM framework for its inflexibility and reliance on functional form assumptions, instead advocating for more generalized, nonparametric formulations of causal mediation effects (Imai et al., 2010a, 2011b; Pearl, 2001, 2014).[3] These are important and valid criticisms that researchers should keep in mind when considering the causal effects they are interested in exploring and estimating. However, the employment of the parametric SEM framework—and, in fact, the use of a linear SEM—is not overly restrictive in the comparative causal mediation context introduced by this paper, as will be explained below.

[3]See Shpitser and VanderWeele (2011) and VanderWeele (2015) for a discussion of the connection between the nonparametric SEM and potential outcomes approaches to causal mediation analysis.

Table 2: Theoretical Implications of Combined Hypotheses

	$\frac{ACME_2}{ACME_1} > 1$	$\frac{\left(\frac{ACME_2}{ATE_2}\right)}{\left(\frac{ACME_1}{ATE_1}\right)} > 1$	
$ATE_2 > ATE_1$	**yes**	**yes**	**Disproportionate scaling up**: Causal channel via M is larger in both absolute and proportional terms for second treatment. M is disproportionately responsible for enhancement of the effect when switching from first to second treatment.
	no	**no**	**Unrelatedness of mediator**: The larger effect of the second treatment is not due to M.
	yes	**no**	**Proportionate scaling up**: Causal channel via M is larger in absolute but not proportional terms for second treatment. M shares responsibility with other causal channels for enhancement of the effect when switching from first to second treatment.
$ATE_2 = ATE_1$	**yes**	**yes**	**Distinct causal anatomies**: Despite equivalent ATEs, the treatments are comprised of differently sized causal channels, with M constituting a larger channel for the second treatment.
	no	**no**	**Indistinguishable causal anatomies**: Any differences in the treatments' causal anatomies are unrelated to M.

Note: Missing yes/no conditions are not applicable.

A. Single Treatment Background

Before proceeding to the comparative causal mediation context, consider the classic linear structural equations model for performing causal mediation (Baron and Kenny, 1986):

$$\begin{aligned} M &= \alpha T + \epsilon_1 \\ Y &= \tau T + \epsilon_2 \\ Y &= \delta T + \beta M + \epsilon_3 \end{aligned}$$

Without loss of generality, mean-centered variables are assumed so as to omit the intercept terms for notational simplicity. T denotes the treatment, M the mediator, and Y the outcome variable. The model implicitly assumes linearity and no interaction between the treatment and mediator. The total effect is τ. The primary effect of interest is the mediation effect, or indirect effect, which is traditionally calculated as the product of α and β. While this product-of-coefficients estimator is specific to the linear SEM, it has been shown to be a special case that nests within more general frameworks of causal mediation (Jo, 2008; Imai et al., 2010b; Pearl, 2014), including the potential outcomes framework. From that perspective, using the

language of the potential outcomes approach, the total effect can also be termed the average treatment effect (ATE), and the indirect effect termed the average causal mediation effect (ACME).[4]

A problem, however, lies in the identification and estimation of the quantity of interest. Specifically, if T has been randomized but M has not, the error terms ϵ_1 and ϵ_3 are likely to be correlated due to unobserved pre-treatment confounders affecting both M and Y, which results in biased and inconsistent estimation of β and hence biased and inconsistent estimation of the ACME. For instance, estimating β by linear regression results in $E[\hat{\beta}] = \beta + \frac{cov(\epsilon_1,\epsilon_3)}{var(\epsilon_1)}$, which includes an asymptotic bias term (Bullock et al., 2008). In other words, unbiased and consistent estimation of the causal mediation effect requires the strong and nonrefutable assumption of no unobserved confounding of the relationship between the mediator and the outcome variable.

This is, in fact, true regardless of the statistical framework used for the identification of causal mediation effects. However, the formal basis of the assumption of no unobserved mediator-outcome confounding takes slightly different forms depending on the statistical framework. In the SEM approach, this takes the form of recursivity or no correlation between the errors of the different equations, while in the potential outcomes framework, the unconfoundedness of the mediator-outcome relationship is implied by the "sequential ignorability" assumption. Notably, methods of sensitivity analysis have been developed to systematically assess the impact of violations of this assumption (e.g. Imai et al., 2010b). However, while such analyses allow for evaluation of the sensitivity of causal mediation estimates, they do not enable the recovery of consistent or unbiased estimates.

Graphically, the problem is depicted by the directed acyclic graph (DAG) in Figure 1. T represents the (exogenous) treatment, Y the outcome variable, M the mediator, and X the collection of unobserved pre-treatment variables that may affect both M and Y. Because X is not fully observed, estimation of the effect of M on Y is biased. Furthermore, estimation of the direct effect of T on Y (i.e. the portion not transmitted via M) is also biased because conditioning on M results in post-treatment bias by virtue of M's status as a collider. In sum, regardless of the statistical framework being employed—whether it is a parametric/nonparametric SEM or the potential outcomes framework—unobserved variables that affect both the mediator and outcome will result in biased and inconsistent estimates of the ACME.

[4]In the potential outcomes approach, the ACME for a single binary treatment is defined formally as $E[Y(t, M(1)) - Y(t, M(0))]$, where the value is the same for $t = 0, 1$ in the case of no interaction between the treatment and mediator.

Figure 1: Confounding in Causal Mediation Analysis

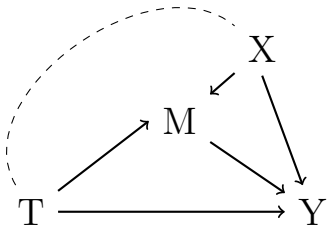

B. Set-up for Comparative Causal Mediation Framework: Two Treatments and Single Mediator

Consider the same model as presented above except that now there are two treatments rather than one.

$$M = \alpha_1 T_1 + \alpha_2 T_2 + \epsilon_1 \quad (1)$$

$$Y = \tau_1 T_1 + \tau_2 T_2 + \epsilon_2 \quad (2)$$

$$Y = \delta_1 T_1 + \delta_2 T_2 + \beta M + \epsilon_3 \quad (3)$$

This set-up is intended for the experimental setting in which there are two mutually exclusive treatment conditions (in addition to a control condition), and thus T_1 and T_2 are binary variables that cannot both be 1 for any single observation. Furthermore, as is common in actual empirical applications of causal mediation, the mediator M is intended to be (though need not be) binary as well. Now, τ_1 and τ_2 correspond to the ATEs for each treatment, the product of α_1 and β corresponds to the ACME for the first treatment with respect to the mediator of interest, and the product of α_2 and β corresponds to the ACME for the second treatment with respect to the mediator of interest. This linear SEM will serve as the framework for the comparative causal mediation approach presented in this paper.[5] Later in the paper, treatment-mediator interactions will be introduced into equation (3).

[5]It is worth explicitly noting that the method presented in this paper does not apply to comparing the effects of a single treatment transmitted via different mediators. In contrast to the method presented in this paper, trying to compare the effects transmitted via multiple mediators would compound the threat to internal validity, as the problem of confounding is likely to affect different mediators to different extents and in ways that cannot be measured or tested. As a separate issue, there is also a possibility of causal connections between the mediators, further threatening clean identification and obscuring what is even being measured or estimated. Guidance on how to handle these issues, which are not covered in this paper, can be found in Imai and Yamamoto (2013) and Daniel et al. (2015).

C. Assumptions

Applied to the intended context and variable characteristics described above, the linear SEM is not as restrictive a model as conventional wisdom would suggest. To further explain, described below are the necessary assumptions for the comparative causal mediation method.[6]

Assumption 1: *Exogeneity of the treatments.*

The exogeneity of the treatments is ensured by random assignment and, hence, is unproblematic in the experimental context.

Assumption 2: *No interaction between the treatments and mediator in equation (3).*

Compared to the first assumption, which can be guaranteed by design, the second assumption is somewhat stringent. However, as will be explained later, it is also testable. Furthermore, this assumption will also be relaxed later.

Assumption 3: *Linear relationships between the outcome, mediator, and treatments.*

The third assumption is that the linear functional form in equations (1)-(3) is correct. While the assumption of linearity may seem rigid and demanding, this assumption is made trivial by the employment of a binary mediator. Given a binary mediator, the two binary treatment indicators, and the no-interaction assumption, the three regression equations are then fully saturated: the two treatments are mutually exclusive, precluding the possibility of $T_1 \cdot T_2$ and $T_1 \cdot T_2 \cdot M$ interactions, while the no-interaction assumption rules out $T_1 \cdot M$ and $T_2 \cdot M$ interactions.[7] As a result, the assumption of linearity is trivial in this context since fully saturated models are "inherently linear" (Angrist and Pischke, 2009a, p. 37), which is why a linear SEM is not an overly restrictive framework in the CCM context. In addition, this means that a binary choice model (e.g. logit, probit) is unnecessary both for equation (1) and, in the case of a binary outcome Y, for equations (2) and (3). Finally, this also helps to justify the exclusion of covariates from the model. In contrast to the case of estimating a single causal mediation effect, the CCM estimands can be estimated consistently without covariate adjustment, as will be shown shortly; furthermore, inclusion of covariates would invalidate the full saturation, and hence linearity, of the model.

[6] As usual, the standard assumption of no interference between units also applies.

[7] The no-interaction assumption will later be relaxed, leading to regression models that are fully saturated without this assumption.

ASSUMPTION 4: $Cov(\alpha_{1i}, \beta_i) = Cov(\alpha_{2i}, \beta_i) = 0$

The assumption that $Cov(\alpha_{1i}, \beta_i) = Cov(\alpha_{2i}, \beta_i) = 0$ is implicit in the standard SEM formulation, whereby the parameters are assumed to be constant structural effects, thereby meaning they do not vary across units and guaranteeing zero covariance across units. The assumption of constant effects is clearly an unrealistic one—and one that is dispensed with in the potential outcomes approach, to its advantage. SEMs also often incorporate varying or heterogeneous coefficients. Applying this approach, the parameters in equations (1)-(3) can be formulated as the mean values of heterogeneous effects, with the effect heterogeneity subsumed into the error terms. While the assumption of constant effects is not necessary for the results that follow, it provides a substantial amount of notational and analytical convenience and hence will be employed for the following results. Later on in this section, subsection G will revisit the assumption of constant effects and show how it is not necessary as long as $Cov(\alpha_{1i}, \beta_i) = Cov(\alpha_{2i}, \beta_i) = 0$, which is an assumption that is also implicitly made in other approaches to causal mediation.[8]

D. Consistency of CCM Estimators

Notably, the method presented here dispenses with the assumption of no confounding of the relationship between the mediator and outcome, which is a strong and nonrefutable assumption that is the most often criticized component of causal mediation analysis (e.g. Gerber and Green, 2012; Bullock et al., 2010; Glynn, 2012; Bullock and Ha, 2011). In fact, with the assumptions above in place, it can be shown that it is still the case that $E[\hat{\beta}] = \beta + \frac{cov(\epsilon_1, \epsilon_3)}{var(\epsilon_1)}$.[9] In contrast, α_j can be estimated consistently and without bias for both $j = 1, 2$. The key implication of these results is that, if comparing two treatments and their mediated effects via the same mediator, then a common bias afflicts both estimated ACMEs. By corollary, this means that the unavoidable mediation bias does not prevent us from comparing the causal mediation anatomies of two different treatments, as long as we are doing so in terms of the same mediator, thereby enabling us to gain important comparative insights on the mediation processes entailed by sets of treatments.

[8]For instance, in the potential outcomes approach as applied to a linear structural form, a conditional version of this assumption is implied by sequential ignorability. As Imai et al. (2010b) note, the sequential ignorability assumption implies a set of assumptions developed by Pearl (2001), which includes the independence between the potential values of the outcome and the potential values of the mediator. In the linear structural form, α_i is a function of the potential values of the mediator, while β_i is a function of the potential values of the outcome. The independence between the potential values of the outcome and the potential values of the mediator implies the independence between these functions, thus implying independence between α_i and β_i.

[9]See Appendix B for proof.

More specifically, it is possible to consistently estimate the two causal mediation comparative estimands of interest:

$$\frac{ACME_2}{ACME_1} = \frac{\alpha_2\beta}{\alpha_1\beta} \qquad \textit{and} \qquad \frac{(\frac{ACME_2}{ATE_2})}{(\frac{ACME_1}{ATE_1})} = \frac{(\frac{\alpha_2\beta}{\tau_2})}{(\frac{\alpha_1\beta}{\tau_1})}$$

THEOREM 1: *Call $\hat{\tau}_2$, $\hat{\tau}_1$, $\hat{\alpha}_2$, $\hat{\alpha}_1$, and $\hat{\beta}$ the linear least squares regression estimators of the parameters from equations (1), (2), and (3). Given no interference between the units and assumptions 1-4, then:*[10]

$$plim\left(\frac{\hat{\alpha}_2\hat{\beta}}{\hat{\alpha}_1\hat{\beta}}\right) = \frac{ACME_2}{ACME_1} \qquad \textit{and} \qquad plim\left(\frac{(\frac{\hat{\alpha}_2\hat{\beta}}{\hat{\tau}_2})}{(\frac{\hat{\alpha}_1\hat{\beta}}{\hat{\tau}_1})}\right) = \frac{(\frac{ACME_2}{ATE_2})}{(\frac{ACME_1}{ATE_1})}$$

In sum, the CCM estimands can be estimated consistently through the simple use of linear least squares regression estimators.

E. Uncertainty Estimation

Because the estimands of interest are nonlinear functions of multiple quantities that must each be estimated, the distributional behavior of the CCM estimators in finite samples is not immediately obvious. Furthermore, because the estimators employ ratios in which the distribution of the denominator may have positive probability density at zero, these estimators do not necessarily have finite-sample moments. This pathological problem is characteristic of ratio estimators in general, and it theoretically complicates the calculation of confidence intervals for those estimators. The existence of probability density at the point where the denominator equals zero creates a singularity in the distribution of a ratio estimator, which can result in the mysterious unbounded confidence interval. Yet traditional methods for constructing confidence sets do not necessarily take this property into account, and it has been shown that "any method which cannot generate unbounded confidence limits for a ratio leads to arbitrary large deviations from the intended confidence level" (von Luxburg and Franz, 2009; Gleser and Hwang, 1987; Koschat, 1987; Hwang, 1995). This issue has been studied extensively, with exact solutions having long ago been derived in special cases—such as a numerator and denominator that are jointly normally distributed (Fieller, 1954)—and approximation techniques based on the bootstrap developed recently for more general cases (Hwang, 1995; von Luxburg and Franz, 2009).

[10]Proofs of theorems can be found in Appendix A.

However, it has also been shown that in spite of the mathematical statistical problems with ratio estimation, the use of standard methods for the practical estimation of confidence intervals yields approximately correct coverage under the reasonable condition that the confidence interval is actually bounded at the desired α level, which is met when the $1 - \alpha$ confidence interval of the denominator does not contain zero (Franz, 2007).[11] These are also the conditions under which the CCM estimators can provide meaningful insights from a conceptual perspective. In those cases where the distribution of the estimated quantity in the denominator is not sufficiently bounded away from zero (i.e. where one of the mediation effects or proportions mediated is approximately zero), then it is not a conceptually meaningful quantity for use in these CCM estimators.[12] In other words, it falls outside of the meaningful scope of comparative causal mediation, a proverbial case of comparing apples to oranges.

In sum, the potentially pathological nature of the CCM estimators is more of a mathematical inconvenience than a serious practical problem, just as is the case with the ratio-based instrumental-variables Wald estimator and the just-identified 2SLS estimator more broadly. As long as the denominator is sufficiently bounded away from zero—i.e. is significantly different from zero at the chosen α level—then standard methods of confidence interval construction, including the Delta Method and standard bootstrap techniques, can be used as a practical matter. Furthermore, in cases in which the denominator is not sufficiently bounded away from zero, the variance and confidence intervals of the estimators would explode, similar to the behavior of the just-identified 2SLS estimator.[13] This protects the researcher from drawing misguided conclusions in such contexts.

Because the CCM estimands are only meaningful in the case that the denominator (for either estimand) is non-zero, it is recommended that a denominator that is statistically significantly different from zero at the chosen α level be treated as a precondition for proceeding with CCM estimation. Not only is this recommended as a conceptual matter, but it also allows the researcher to avoid the problems of unbounded confidence intervals and to use the Delta Method and/or standard bootstrap techniques to compute confidence intervals.

[11] And, as in any situation, a sufficiently large sample size is also necessary for analytic methods that rely on the CLT and for bootstrap methods, which rely on the sample distribution adequately approximating the population distribution.

[12] The analog in the instrumental variables estimation is a weak first-stage effect.

[13] For discussion of the finite-sample properties of just-identified 2SLS, including its lack of finite-sample moments and its unstable variance behavior given a weak first stage, see Angrist and Pischke (2009a, pp. 205-218), Angrist and Pischke (2009b), and Nelson and Starz (1988).

F. Finite-Sample Adjustments

Due to their ratio form, the CCM estimators are not exactly centered on the true estimand in finite samples. As the sample size grows, this divergence becomes negligible. In smaller samples, finite-sample adjustments can be derived using Taylor series expansion. This results in the following adjusted estimator for the first CCM estimand:

$$\frac{\hat{\alpha}_2}{\hat{\alpha}_1} + \frac{\widehat{Cov}(\hat{\alpha}_1, \hat{\alpha}_2)}{\hat{\alpha}_1^2} - \frac{\widehat{Var}(\hat{\alpha}_1)\hat{\alpha}_2}{\hat{\alpha}_1^3}$$

Similarly, the adjusted estimator for the second CCM estimand is the following:

$$\begin{aligned}&\frac{\hat{\alpha}_2\hat{\tau}_1}{\hat{\alpha}_1\hat{\tau}_2} - \widehat{Var}(\hat{\alpha}_1)\frac{\hat{\alpha}_2\hat{\tau}_1}{\hat{\alpha}_1^3\hat{\tau}_2} - \widehat{Var}(\hat{\tau}_2)\frac{\hat{\alpha}_2\hat{\tau}_1}{\hat{\alpha}_1\hat{\tau}_2^3} + \widehat{Cov}(\hat{\alpha}_2, \hat{\alpha}_1)\frac{\hat{\tau}_1}{\hat{\alpha}_1^2\hat{\tau}_2} + \widehat{Cov}(\hat{\alpha}_2, \hat{\tau}_2)\frac{\hat{\tau}_1}{\hat{\alpha}_1\hat{\tau}_2^2} \\ &-\widehat{Cov}(\hat{\alpha}_2, \hat{\tau}_1)\frac{1}{\hat{\alpha}_1\hat{\tau}_2} - \widehat{Cov}(\hat{\alpha}_1, \hat{\tau}_2)\frac{\hat{\alpha}_2\hat{\tau}_1}{\hat{\alpha}_1^2\hat{\tau}_2^2} + \widehat{Cov}(\hat{\alpha}_1, \hat{\tau}_1)\frac{\hat{\alpha}_2}{\hat{\alpha}_1^2\hat{\tau}_2} + \widehat{Cov}(\hat{\tau}_2, \hat{\tau}_1)\frac{\hat{\alpha}_2}{\hat{\alpha}_1\hat{\tau}_2^2}\end{aligned}$$

Appendix C presents these results in detail. Simulations, presented in a later section, illustrate the improvement of the adjusted estimators over the simple estimators in small samples.

G. Revisiting the Assumption of Constant Parameters

Consider again the linear structural equations model presented by equations (1)-(3), but now allowing for varying/heterogeneous coefficients:

$$\begin{aligned}M &= \alpha_{1i}T_1 + \alpha_{2i}T_2 + \epsilon_{1i} \\ Y &= \tau_{1i}T_1 + \tau_{2i}T_2 + \epsilon_{2i} \\ Y &= \delta_{1i}T_1 + \delta_{2i}T_2 + \beta_i M + \epsilon_{3i}\end{aligned}$$

Following Imai and Yamamoto (2013), the varying coefficients can be decomposed into their means and the deviations from their means, such that we have:

$$M = \alpha_1 T_1 + \alpha_2 T_2 + \eta_{1i} \quad (4)$$

$$Y = \tau_1 T_1 + \tau_2 T_2 + \eta_{2i} \quad (5)$$

$$Y = \delta_1 T_1 + \delta_2 T_2 + \beta M + \eta_{3i} \quad (6)$$

where the parameters in these equations represent the means of the varying coefficients (e.g. $\alpha_1 = E[\alpha_{1i}]$) and the deviations of the coefficients from their means have been subsumed into

the error terms η_{ji}, such that

$$\begin{aligned}\eta_{1i} &= \tilde{\alpha}_{1i}T_1 + \tilde{\alpha}_{2i}T_2 + \epsilon_{1i} \\ \eta_{2i} &= \tilde{\tau}_{1i}T_1 + \tilde{\tau}_{2i}T_2 + \epsilon_{2i} \\ \eta_{3i} &= \tilde{\delta}_{1i}T_1 + \tilde{\delta}_{2i}T_2 + \tilde{\beta}_i M + \epsilon_{3i}\end{aligned}$$

While this model allows for varying coefficients, as before, estimation of equations (4)-(6) by linear regression will again recover consistent and unbiased estimates of α_1, α_2, τ_1, and τ_2, where these parameters are now mean effects rather than constant effects. Also as before, the estimate of β, now a mean effect, will be biased and inconsistent because $Cov(\epsilon_{1i}, \epsilon_{3i}) \neq 0$ implies $Cov(\eta_{1i}, \eta_{3i}) \neq 0$. However, also as before, the ratio form of the CCM estimators, $\frac{\hat{\alpha}_2\hat{\beta}}{\hat{\alpha}_1\hat{\beta}}$ and $\frac{\left(\frac{\hat{\alpha}_2\hat{\beta}}{\hat{\tau}_2}\right)}{\left(\frac{\hat{\alpha}_1\hat{\beta}}{\hat{\tau}_1}\right)}$, eliminates the bias in the $\hat{\beta}$ terms.

However, in this case, $ACME_j = E[\alpha_{ji}\beta_i]$, whereas the product of estimated coefficients $\hat{\alpha}_j\hat{\beta} = \widehat{E[\alpha_{ji}]}\widehat{E[\beta_i]}$. Recalling that $\hat{\beta} = \widehat{E[\beta_i]}$ is asymptotically biased but the ratio form of the estimands eliminates the asymptotic bias, we can see that $plim\ \frac{\widehat{E[\alpha_{2i}]}\widehat{E[\beta_i]}}{\widehat{E[\alpha_{1i}]}\widehat{E[\beta_i]}} = \frac{E[\alpha_{2i}]E[\beta_i]}{E[\alpha_{1i}]E[\beta_i]}$ and $plim\ \frac{\left(\frac{\widehat{E[\alpha_{2i}]}\widehat{E[\beta_i]}}{\widehat{E[\tau_{2i}]}}\right)}{\left(\frac{\widehat{E[\alpha_{1i}]}\widehat{E[\beta_i]}}{\widehat{E[\tau_{1i}]}}\right)} = \frac{\left(\frac{E[\alpha_{2i}]E[\beta_i]}{E[\tau_{2i}]}\right)}{\left(\frac{E[\alpha_{1i}]E[\beta_i]}{E[\tau_{1i}]}\right)}$. If $Cov(\alpha_{ji}, \beta_i) = 0$, then $E[\alpha_{ji}]E[\beta_i] = E[\alpha_{ji}\beta_i] = ACME_j$, in which case the estimators presented above remain consistent for both estimands.

In sum, for all of the results presented above, the assumption of constant coefficients is not necessary, as long as it is assumed that $Cov(\alpha_{ji}, \beta_i) = 0$ for $j = 1, 2$.

V. Simulations

To illustrate the properties of the CCM method, this section presents a simulation. Simulated causal mediation data were generated according to the following model, with the output of the first equation, M, feeding into the second equation:

$$\begin{aligned}M_i &= \alpha_{1i}T_{1i} + \alpha_{2i}T_{2i} + \psi_i X_i + \nu_{1i} \\ Y_i &= \delta_{1i}T_{1i} + \delta_{2i}T_{2i} + \beta_i M_i + \phi_i X_i + \nu_{3i}\end{aligned}$$

T_1 and T_2 are indicator variable vectors and were generated such that an equal number of units were randomly assigned to (a) neither treatment, (b) T_1, and (c) T_2, with no units assigned to both T_1 and T_2. The rest of the variables and parameters were generated as follows:

$$X \sim Unif(0,5) \quad \alpha_1 \sim N(4,2) \quad \alpha_2 \sim N(10,2) \quad \beta \sim N(3,2)$$
$$\delta_1 \sim N(5,2) \quad \delta_2 \sim N(5,2) \quad \psi \sim N(4,2) \quad \phi \sim N(4,2) \quad \nu_1 \sim N(0,1) \quad \nu_3 \sim N(0,1)$$

As indicated, the parameters were generated to vary independently across units, yielding heterogeneous effects with zero covariance between α_j and β for $j = 1, 2$. Further, the data were also generated with no interaction between T_j and M for $j = 1, 2$. Along with the linear form and the exogeneity of T_j for $j = 1, 2$, all assumptions established above are met by the data-generating process.

Once the data were generated, the mean values of the parameters α_1, α_2, and β—as well as τ_1 and τ_2—were estimated according to equations (4)-(6). Thus X was omitted from the estimation process, simulating unobserved confounders leading to non-zero covariance between the errors in equations (4) and (6), η_1 and η_3, hence also biasing the estimation of β and the $ACMEs$.

In the results presented in Figure 2, the model was simulated 100 times with a total of 300 units per simulation (100 assigned to each of the two treatments and 100 assigned to neither treatment). Each panel in the plot displays the point estimates from each simulation for a different estimand, along with 95% confidence intervals constructed via the nonparametric pairs bootstrap. The panels in the top row correspond to the traditional causal mediation estimands: $ACME_1$ ($\alpha_1\beta$), $ACME_2$ ($\alpha_2\beta$), proportion of T_1 mediated ($\frac{\alpha_1\beta}{\tau_1}$), and proportion of T_2 mediated ($\frac{\alpha_2\beta}{\tau_2}$). The panels in the bottom row correspond to the CCM estimands, with both simple and small-sample adjusted estimators presented. The panels note the coverage of the confidence intervals, the true value of the estimand according to the data-generating process, and the mean estimate over all 100 simulations.

As can be seen, Figure 2 clearly shows how the traditional $ACME$ estimators (top row) are biased and exhibit confidence-interval under-coverage given the presence of unmeasured confounders (X). The top left two panels show that the estimators of $ACME_1$ and $ACME_2$ are biased upward by approximately 2.5 and 6, resulting in only 90% and 72% coverage of the 95% confidence intervals. The story is the same for the top right two panels, which show the estimates of the proportions mediated for each treatment.

In contrast to the clear bias—and inconsistency, though this cannot be seen in this single set of simulations with fixed n—of the traditional causal mediation estimators, the bottom row shows that the CCM estimators are properly centered and exhibit good coverage. The bottom left two panels present the estimators of the ACME ratio, the first being the simple estimator and the second being the small-sample adjusted estimator. As can be seen, both perform well in recovering a mean estimate close to the true estimand value and proper confidence interval coverage (subject to simulation error). In addition, the small-sample adjustment also appears to yield a better mean estimate, demonstrating the gains from this adjustment. The results

Figure 2: Comparative Causal Mediation Simulation: Without Interactions

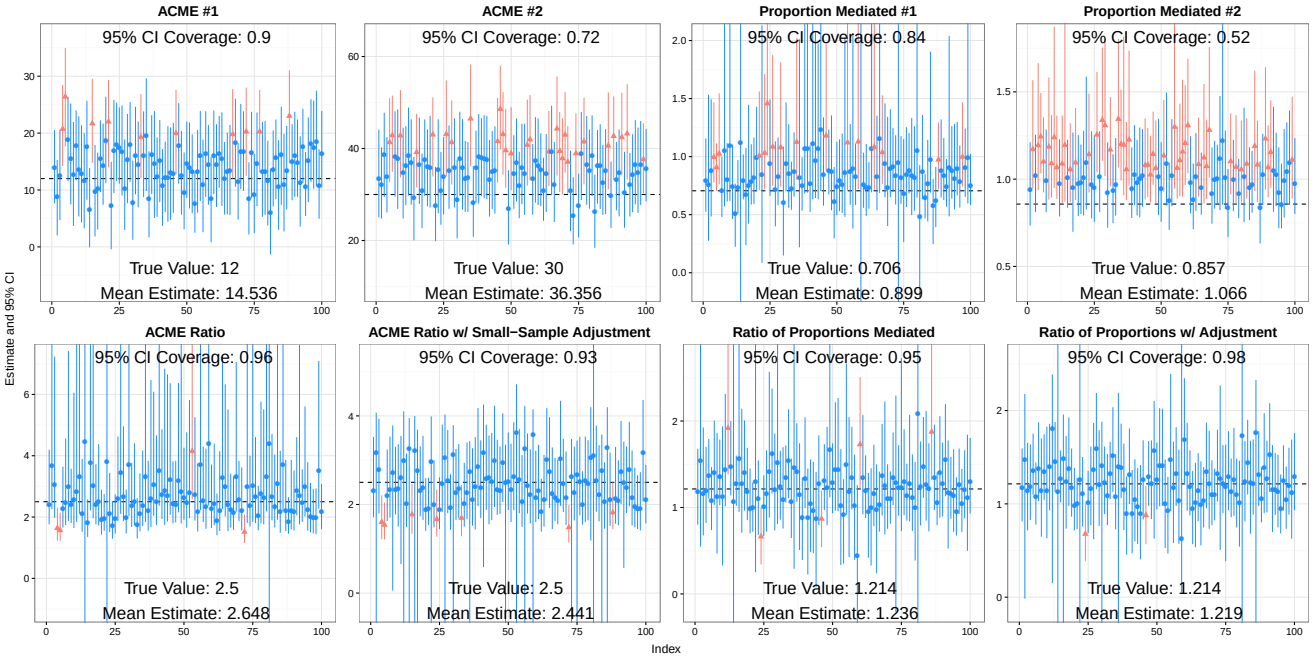

are the same in the bottom right two panels, which show the simple and adjusted estimators for the ratio of proportions mediated.

VI. Relaxing the No-Interaction Assumption

A. *Set-Up*

Consider the same model of linear structural equations but now allowing for treatment-mediator interactions, which has been referred to by some scholars as a version of moderated mediation (James and Brett, 1984; Preacher, 2007). The model, again assuming demeaned variables so as to omit the intercepts, becomes the following:

$$M = \alpha_1 T_1 + \alpha_2 T_2 + \epsilon_1 \tag{7}$$

$$Y = \tau_1 T_1 + \tau_2 T_2 + \epsilon_2 \tag{8}$$

$$Y = \delta_1 T_1 + \delta_2 T_2 + \beta M + \gamma_1 T_1 M + \gamma_2 T_2 M + \epsilon_3 \tag{9}$$

In this case, of interest are functions of the ACMEs for subsamples, namely the treated (ACMETs) and the control (ACMECs). Following Imai et al. (2010a), these can be expressed

as:

$$\begin{aligned} ACMET_1 &= \alpha_1(\beta + \gamma_1) = \alpha_1\omega_1 \\ ACMET_2 &= \alpha_2(\beta + \gamma_2) = \alpha_2\omega_2 \\ ACMEC_1 &= \alpha_1\beta \\ ACMEC_2 &= \alpha_2\beta \end{aligned}$$

The same results as presented above continue to apply in this case with regards to the ACMECs. However, the CCM estimands are likely to be of greater theoretical and practical interest in terms of the ACMETs. In this case, the estimands of interest are as follows:

DEFINITION 2: *To allow for interactions between the treatments and mediator, define the estimands of interest as follows:*

$$Estimand\,1: \frac{ACMET_2}{ACMET_1}$$

$$Estimand\,2: \frac{(\frac{ACMET_2}{ATE_2})}{(\frac{ACMET_1}{ATE_1})}$$

B. Conservatism of Estimators

Call $\hat{\tau}_2$, $\hat{\tau}_1$, $\hat{\alpha}_2$, $\hat{\alpha}_1$, $\hat{\beta}$, $\hat{\gamma}_2$, and $\hat{\gamma}_1$ the linear least squares regression estimators of the parameters from equations (7), (8), and (9). Once again, the exogeneity of the treatments guarantees unbiasedness for $\hat{\tau}_2$, $\hat{\tau}_1$, $\hat{\alpha}_2$, and $\hat{\alpha}_1$, but not for $\hat{\beta}$, $\hat{\gamma}_2$, and $\hat{\gamma}_1$.[14] Under certain conditions, it can be shown that

$$\frac{\hat{\alpha}_2(\hat{\beta}+\hat{\gamma}_2)}{\hat{\alpha}_1(\hat{\beta}+\hat{\gamma}_1)} = \frac{\hat{\alpha}_2\hat{\omega}_2}{\hat{\alpha}_1\hat{\omega}_1} \qquad \textit{and} \qquad \frac{\left(\frac{\hat{\alpha}_2(\hat{\beta}+\hat{\gamma}_2)}{\hat{\tau}_2}\right)}{\left(\frac{\hat{\alpha}_1(\hat{\beta}+\hat{\gamma}_1)}{\hat{\tau}_1}\right)} = \frac{\left(\frac{\hat{\alpha}_2\hat{\omega}_2}{\hat{\tau}_2}\right)}{\left(\frac{\hat{\alpha}_1\hat{\omega}_1}{\hat{\tau}_1}\right)}$$

are not consistent estimators of $\frac{ACMET_2}{ACMET_1}$ and $\frac{(\frac{ACMET_2}{ATE_2})}{(\frac{ACMET_1}{ATE_1})}$, respectively, but are consistently conservative (attenuated toward unity) under reasonable and testable conditions. These simple estimators are conservative only in the probability limit because, as before, there is a finite-sample divergence due to the ratio form of the estimators. However, also as before, that finite-sample divergence can be approximated, estimated, and used to construct adjusted estimators that are not just conservative in the probability limit but also conservative in small samples.

[14] Loeys et al. (2016) describe specific conditions under which $\hat{\gamma}_2$ and $\hat{\gamma}_1$ are unbiased estimators even when $\hat{\beta}$ is not.

THEOREM 2: *Without loss of generality, assume that both the numerator and denominator of the estimand of interest are positive, and that the estimand is greater than 1 (i.e. the numerator is larger than the denominator). Call* $\hat{\tau}_2, \hat{\tau}_1, \hat{\alpha}_2, \hat{\alpha}_1, \hat{\beta}, \hat{\gamma}_2, \hat{\gamma}_1$ *the linear least squares regression estimators of the parameters from equations (7), (8), and (9). Let* $\hat{\omega}_1 = \hat{\beta} + \hat{\gamma}_1$ *and* $\hat{\omega}_2 = \hat{\beta} + \hat{\gamma}_2$. *Further call* ξ_1 *and* ξ_2 *the bias components of* $\hat{\omega}_1$ *and* $\hat{\omega}_2$, *respectively (i.e.* $E[\hat{\omega}_1] - \omega_1 = \xi_1$ *and* $E[\hat{\omega}_2] - \omega_2 = \xi_2$*). In addition to no interference between units, make assumptions 1 and 3, as well as a revised version of assumption 4:* $Cov(\alpha_{1i}, \omega_{1i}) = Cov(\alpha_{2i}, \omega_{2i}) = 0$. *Then, given* $\xi_1 > 0$, $\xi_2 > 0$, *and* $\omega_2\xi_1 > \omega_1\xi_2$, *the following holds:*

$$plim\ \frac{\hat{\alpha}_2\hat{\omega}_2}{\hat{\alpha}_1\hat{\omega}_1} < \frac{ACMET_2}{ACMET_1} \qquad and \qquad plim\ \frac{\left(\frac{\hat{\alpha}_2\hat{\omega}_2}{\hat{\tau}_2}\right)}{\left(\frac{\hat{\alpha}_1\hat{\omega}_1}{\hat{\tau}_1}\right)} < \frac{\left(\frac{ACMET_2}{ATE_2}\right)}{\left(\frac{ACMET_1}{ATE_1}\right)}$$

The result is that, given the conditions described in Theorem 2, the bias involved in estimating $\frac{ACMET_2}{ACMET_1}$ and $\frac{\left(\frac{ACMET_2}{ATE_2}\right)}{\left(\frac{ACMET_1}{ATE_1}\right)}$ results in attenuated estimates of these estimands. Since these results were presented without loss of generality in the context where the estimand is greater than 1, this means that the attentuated estimates will be conservative. In other words, the estimates will be biased in favor of the null hypothesis that the estimand equals 1.

While assumption 2 (no interaction between the treatments and mediator) was relaxed, Theorem 2 introduces the folllowing additional conditions that were not present in Theorem 1: $\xi_1 > 0$, $\xi_2 > 0$, and $\omega_2\xi_1 > \omega_1\xi_2$. These two conditions can be assessed empirically, as discussed below.

C. Testing the Assumption on the Sign of ξ_j

It is worth noting that in the most common case of "consistent mediation"—which is where the total treatment effects and mediation effects are of the same sign, in contrast to "inconsistent mediation" where the mediator actually suppresses the treatment effects (MacKinnon et al., 2007)—the bias of an estimated mediation effect may be expected to have the same sign as the true mediation effect, as such biases may result from the theoretically likely case of pretreatment confounders affecting the mediator and outcome in the same direction. As stated in Theorem 2, the assumption is made without loss of generality that the mediation effects are positive; thus, we would expect the bias terms also to be positive.

That being said, it is useful to be able to test this latter expectation. We know that ξ_1 and ξ_2 will have a common sign, since they are ratios that share a common numerator and both have positive denominators. Furthermore, the assumption that ξ_1 and ξ_2 are of a particular

sign can be indirectly tested and diagnosed. This requires that important covariates believed to be related to both the outcome and mediator variables are observed. The inclusion of these covariates in the estimation of each $ACMET$ can then be investigated. If, as additional covariates are included in equation (9) of the estimation model, the $ACMET$ estimates get attenuated toward zero, this is evidence that their omission biases the estimates in the direction of the sign of the mediation effects. If it can be argued that these covariates encompass many or all of the most important confounding variables (i.e. it is unlikely that the omission of some other unobserved confounder is biasing the estimates on net in the other direction, toward zero), this provides plausible evidence that the overall bias ξ_j shares the same sign as the mediated effects in the uncontrolled models that are used in the actual estimation process.

An illustration of this test is included later in subsection D of section VII, which presents an application of the CCM method.

D. Testing the Assumption that $\omega_2\xi_1 > \omega_1\xi_2$

The condition $\omega_2\xi_1 > \omega_1\xi_2$ implies that:

$$\omega_2 Var(\epsilon_1^{T_2}) > \omega_1 Var(\epsilon_1^{T_1})$$

where $\epsilon_1^{T_1}$ ($\epsilon_1^{T_2}$) denotes the error term in equation (7) exclusively for the subset that has been assigned to the first (second) treatment condition.[15] Since, without loss of generality, we are considering the hypothesis that $\frac{ACMET_2}{ACMET_1} > 1$ (or $\frac{\left(\frac{ACMET_2}{ATE_2}\right)}{\left(\frac{ACMET_1}{ATE_1}\right)} > 1$), it would also be reasonable to expect that the interaction between T_j and M is larger for T_2 than T_1 (i.e. $\gamma_2 > \gamma_1$), in which case $\omega_2 > \omega_1$.

Further, the assumption that $\omega_2 > \omega_1$ is testable by comparing $\hat{\omega}_2$ and $\hat{\omega}_1$. In fact, as shown in the appendix, $\omega_2 Var(\epsilon_1^{T_2}) > \omega_1 Var(\epsilon_1^{T_1})$ implies asymptotically (and in expectation) that $\hat{\omega}_2 Var(\epsilon_1^{T_2}) > \hat{\omega}_1 Var(\epsilon_1^{T_1})$. In addition, $Var(\epsilon_1^{T_1})$ and $Var(\epsilon_1^{T_2})$ can also be estimated simply by computing the variance of M for each treatment condition. Thus, the assumption that $\omega_2\xi_1 > \omega_1\xi_2$ is directly testable by the following:

$$\hat{\omega}_2 \widehat{Var}(\epsilon_1^{T_2}) > \hat{\omega}_1 \widehat{Var}(\epsilon_1^{T_1})$$

E. Sensitivity Analysis

To review, the no-interaction assumption allows for establishing the consistency of the CCM estimators, as presented in section IV.D. In contrast, when the no-interaction assumption is

[15] See proof of Theorem 2 in Appendix A for more details.

relaxed, the CCM estimators are no longer consistent. However, as has been shown in this section (section VI), the CCM estimators are asymptotically conservative under additional testable conditions.

In particular, the conservatism of the CCM estimators depends upon the testable assumption that $\omega_2\xi_1 > \omega_1\xi_2$. In situations where the researcher has reason to doubt that assumption—for instance, if $\hat{\omega}_2\widehat{Var}(\epsilon_1^{T_2}) < \hat{\omega}_1\widehat{Var}(\epsilon_1^{T_1})$ because $\hat{\omega}_2 < \hat{\omega}_1$—a sensitivity analysis can be performed to understand how sensitive the estimates are to the unobserved bias. To do so, the sources of bias can be conceptualized as a single composite confounding variable that is scaled to the interval $[0, 1]$, which will be called X. As shown in the appendix, the asymptotic bias on the estimators is $\frac{\omega_1(\omega_2+\xi_2)}{\omega_2(\omega_1+\xi_1)}$, which can be re-expressed as:

$$\frac{\omega_1(\omega_2 + \xi_2)}{\omega_2(\omega_1 + \xi_1)} = plim\, \frac{\hat{\omega}_2\left(\hat{\omega}_1 - \frac{\psi\phi Var(X)}{Var(\epsilon_1^{T_1})}\right)}{\hat{\omega}_1\left(\hat{\omega}_2 - \frac{\psi\phi Var(X)}{Var(\epsilon_1^{T_2})}\right)}$$

The terms $\hat{\omega}_1$ and $\hat{\omega}_2$ are observed since they are estimated, and as explained above $Var(\epsilon_1^{T_1})$ and $Var(\epsilon_1^{T_2})$ can also be estimated. $Var(X)$ can be set at its maximum possible value of 0.5, which is conservative (i.e. maximizes the bias term) in the case that $\hat{\omega}_2 < \hat{\omega}_1$. Finally, the terms ψ and ϕ denote the effects of the composite confounder X on M and Y, respectively, given a change in the value of X from its minimum (0) to its maximum (1). The researcher can plug in hypothetical values for ψ and ϕ, based on subject-matter expertise on the likely magnitude of confounding present in the subject under study, in order to compute hypothetical bias values. The estimates of $\frac{ACMET_2}{ACMET_1}$ and $\frac{\left(\frac{ACMET_2}{ATE_2}\right)}{\left(\frac{ACMET_1}{ATE_1}\right)}$ can then be multiplied by the inverse of the hypothetical bias values to yield implied true values of $\frac{ACMET_2}{ACMET_1}$ and $\frac{\left(\frac{ACMET_2}{ATE_2}\right)}{\left(\frac{ACMET_1}{ATE_1}\right)}$. These implied values can then be compared to the respective estimates to assess the sensitivity of the estimators to various degrees of bias (i.e. various values of ψ and ϕ).

F. Additional Notes

Similar to the case in which the no-interaction assumption is maintained, finite-sample adjustments can be derived for the CCM estimators when relaxing the no-interaction assumption. Appendix C presents these finite-sample adjustments. In addition, Appendix D presents simulation results to illustrate the properties of the CCM estimators when the no-interaction assumption has been relaxed.

VII. Application

A. Background

Does international law affect state behavior? There is a longstanding scholarly debate on this question, with some political scientists and legal scholars viewing international law as epiphenomenal to state interests and power, and others seeing international law as having a real impact on state decision-making. Among the latter group, many scholars have identified domestic political processes and institutions as an important conduit through which national governments can be induced to honor their international legal obligations, even in cases where those governments did not intend to comply in the first place. The electoral compliance mechanism, in which governments are incentivized to maintain compliance with international legal agreements under the threat of electoral punishment for violations, is one such domestic source of compliance that has been theorized. In a number of recent studies using survey experiments, political scientists have accumulated evidence that voters in the United States and elsewhere are inclined to punish elected officials who renege on previous foreign policy commitments (Tomz, 2007; Hyde, 2015). These electoral costs for breaking commitments are known as domestic "audience costs."

However, an important gap remains in this scholarship: while studies have shown that public disapproval of a foreign policy decision tends to increase when that policy decision requires reneging on international legal commitments, these studies have not isolated the role of legality per se in generating that disapproval. In other words, the design of these studies has masked the extent to which such disapproval is attributable to the baseline breaking of the commitment (i.e. the audience costs for not honoring a policy pledge in general) versus the additional legal status of the commitment.

B. Study Design

In order to address this gap in the literature, the author designed and implemented a novel survey experiment embedded in an online survey administered in August 2015, with 1602 respondents recruited via Amazon Mechanical Turk. The experiment revolved around a security scenario in which the U.S. government decided to take military action against ISIS forces in Iraq.[16] The survey instrument text can be found in Appendix E, and basic sample demographic

[16]The scenario involved a U.S. military operation in Iraq to capture ISIS militants who were threatening rocket attacks on neighboring countries but were hiding in a civilian zone. Respondents were told that in order to avoid collateral damage, the U.S. deployed commandos in a covert operation, in which the commandos used an ostensibly non-lethal incapacitating chemical gas to neutralize the ISIS militants. The incapacitating gas was featured in the scenario in order to exploit real-world ambiguity (and contention)

distributions can be found in Appendix F.

There were two primary goals of the research. The first goal was to disentangle the dimension of (il)legality from the baseline violation of a foreign policy commitment more explicitly than have previous studies, thereby creating a more valid design to answer the key research question: Does the international legal status of a foreign policy commitment increase the potential for domestic audience costs if that commitment is violated? To achieve this goal, the experimental design featured two mutually exclusive treatment conditions in addition to a control condition. In the control condition, respondents were simply told about the U.S. government's decision to use military force. In the first "non-legal" treatment condition, respondents were additionally told that this decision constituted a violation of the U.S. government's previous foreign policy commitment. In the second "legal" treatment condition, respondents were told that this decision constituted a violation of the U.S. government's international legal commitment. There were two binary outcome variables of interest. The first measured whether or not the respondent disapproved of the policy decision, which will be called Disapproval. The second measured whether or not the respondent would be less likely to vote for a U.S. Senator who supported the policy decision, which will be called Punishment.

The second research goal was to identify and better understand the contours of public opinion that determine the extent to which legalization does (or does not) amplify audience costs. In addition to measuring Disapproval and Punishment, certain attitudes were also measured for investigation as mediators. The primary mediator of interest was the respondents' perception of the (im)morality of the decision to use military force, which represents a key theoretical reason, long noted by scholars, that voters would disapprove more greatly of violations of legalized foreign policy commitments than similar non-legalized commitments (Abbott and Snidal, 2000, pp. 428-429). Hence, a binary mediator variable measured whether or not each respondent believed the policy decision to be immoral, which will be called Perceived Immorality. This enables estimation of the portion of each treatment effect, ATE_1 (non-legal) and ATE_2 (legal), that is transmitted via Perceived Immorality—that is, estimation of $ACME_1$ and $ACME_2$.

surrounding the international legality of chemical incapacitants in unconventional operations, as well as ambiguity surrounding the lethality of these chemical agents. Because of this ambiguity and the technical nature of the legal categorization of chemical incapacitants, average voters should not be expected to identify such agents as clearly illegal, in contrast to well-known chemical warfare agents. At the same time, that these chemical incapacitants are illegal under the Chemical Weapons Convention, which is probably the most widely accepted position among arms control legal experts (though some argue otherwise in terms of the use of chemical incapacitants under certain conditions; for an overview of the debate, see Ballard (2007)) is also plausible, enabling clean and easy administration of the broken legal commitment treatment.

As described above, the problem with traditional mediation analysis is that, even with pre-treatment covariates included as controls, those mediated effects are likely to be biased and inconsistent. However, under the assumptions stated earlier, the CCM estimands can be estimated consistently (or conservatively). The first estimand $\frac{ACME(T)_2}{ACME(T)_1}$ measures the extent to which the morality mediator transmits a stronger effect for the legal treatment than for the non-legal treatment. The second estimand $(\frac{ACME(T)_2}{ATE_2})/(\frac{ACME(T)_1}{ATE_1})$ measures the extent to which the morality mediator comprises a larger proportion of the total effect of (i.e. is more important for) the legal treatment, compared to the non-legal treatment.

C. Results

The results of the survey experiment provided statistically and substantively strong evidence that the legal treatment did indeed cause a larger increase in the probability of Disapproval and Punishment than the non-legal treatment, as shown by Table 3, providing support for the theory that legalization enhances audience costs. Specifically, the legal treatment had an estimated 12.5 percentage-point larger effect on the probability of Disapproval and a 9.9 percentage-point larger effect on the probability of Punishment than the non-legal treatment.

Table 3: Sample Estimates of ATEs

	DV: Disapproval		
	$\widehat{ATE}_1$ (Non-legal treatment effect)	$\widehat{ATE}_2$ (Legal treatment effect)	$\widehat{ATE}_2 - \widehat{ATE}_1$ (Difference in treatment effects)
Estimate	0.195	0.320	0.125
Bootstrap 95% CI	[0.134, 0.249]	[0.265, 0.372]	[0.065, 0.181]
	DV: Punishment		
	$\widehat{ATE}_1$ (Non-legal treatment effect)	$\widehat{ATE}_2$ (Legal treatment effect)	$\widehat{ATE}_2 - \widehat{ATE}_1$ (Difference in treatment effects)
Estimate	0.182	0.281	0.099
Bootstrap 95% CI	[0.127, 0.234]	[0.226, 0.333]	[0.039, 0.156]

More importantly in the context of this paper, however, the results of the CCM analysis also provided support for the theory that this enhancement of audience costs by legalization

is, at least in part, due to an increase in Perceived Immorality. Table 4 shows the results of the CCM analysis. The assumption of no interaction between the treatments and mediator was tested in the case of both dependent variables. The test failed to reject the null hypothesis of no interactions in the case of the Disapproval dependent variable, and hence the no-interaction assumption was maintained in that case.

However, the test rejected the null hypothesis of no interactions in the case of the Punishment dependent variable, which is why the causal mediation estimates in the Punishment case involve the ACMETs (ACMEs for the treated). Furthermore, additional tests provide support for the conditions necessary for the CCM estimators to be conservative given the interactions between the treatments and mediator. Specifically, estimates revealed indeed that $\hat{\omega}_2\widehat{Var}(\epsilon_1^{T_2}) > \hat{\omega}_1\widehat{Var}(\epsilon_1^{T_1})$, while diagnostics also provided evidence for a positive asymptotic bias, as described in subsection D below.

Table 4: Comparative Causal Mediation Analysis

	$\frac{\widehat{ACME}_2}{\widehat{ACME}_1}$	$\frac{\left(\frac{\widehat{ACME}_2}{\widehat{ATE}_2}\right)}{\left(\frac{\widehat{ACME}_1}{\widehat{ATE}_1}\right)}$	$\widehat{ACME}_1$	$\widehat{ACME}_2$
DV: Disapproval				
Estimate	**1.563**	**0.952**	0.113	0.177
Bootstrap 95% CI	**[1.189, 2.170]**	**[0.752, 1.212]**	[0.078, 0.149]	[0.140, 0.215]
	$\frac{\widehat{ACMET}_2}{\widehat{ACMET}_1}$	$\frac{\left(\frac{\widehat{ACMET}_2}{\widehat{ATE}_2}\right)}{\left(\frac{\widehat{ACMET}_1}{\widehat{ATE}_1}\right)}$	$\widehat{ACMET}_1$	$\widehat{ACMET}_2$
DV: Punishment				
Estimate	**1.829**	**1.184**	0.096	0.176
Bootstrap 95% CI	**[1.321, 2.657]**	**[0.886, 1.635]**	[0.064, 0.131]	[0.137, 0.213]

The CCM estimates are presented in bold in Table 4. Given the large sample size, these estimates were obtained using the simple estimators,[17] and the 95% confidence intervals were

[17]The finite-sample adjusted estimates are virtually identical, as should be expected given the sample size. For instance, the adjusted estimate of $\frac{\widehat{ACME}_2}{\widehat{ACME}_1}$ for the Disapproval dependent variable is 1.533, and the adjusted estimate of $\frac{\widehat{ACMET}_2}{\widehat{ACMET}_1}$ for the Punishment dependent variable is 1.796.

computed via the nonparametric pairs bootstrap. As can be seen, the $\frac{\widehat{ACME(T)_2}}{\widehat{ACME(T)_1}}$ estimates are statistically (and substantively) distinguishable from 1 for both dependent variables. These estimates can be interpreted as meaning that the effect on Disapproval (Punishment) mediated via Perceived Immorality is about 56% (83%) larger for the legal treatment than for the non-legal treatment. In contrast, the $\frac{\left(\frac{\widehat{ACME(T)_2}}{\widehat{ATE}_2}\right)}{\left(\frac{\widehat{ACME(T)_1}}{\widehat{ATE}_1}\right)}$ estimates are not statistically distinguishable from 1 for either dependent variable. This means that while Perceived Immorality transmitted a larger effect for the legal treatment than the non-legal treatment, it did not constitute a larger proportion of the overall ATE for the legal treatment compared to the non-legal treatment.

In combination, the results correspond to the case of "proportionate scaling up" presented in Table 2. Perceived Immorality (M) is found to be an important factor that leads to a scaling up of the audience costs effect given legalization. Yet it also appears that other mediation channels also help to scale up that effect such that while the mediation channel via Perceived Immorality expands, it does not increase as a proportion of the total effect. Overall, the collective evidence provides support in favor of the theory that international legalization enhances audience costs (at least partially) by amplifying the perceived immorality of violating a foreign policy commitment.

As explained in this paper, it is difficult to have confidence in the accuracy of the individual $\widehat{ACME(T)}$ estimates, even if covariates were to be included as controls, because this would require a strong assumption of no unobserved confounding of the mediator-outcome relationship. However, under a weaker set of assumptions, one can have greater confidence in the estimates of the CCM estimands, which also happen to be the estimands of primary interest for the particular theory and hypotheses at stake in this and many other applications.

D. Illustration of Test of the Sign of the Bias

In the application presented above, the no-interaction assumption is tested and rejected in the case of the Punishment dependent variable. As noted earlier, in a case in which the no-interaction assumption is relaxed, additional assumptions must be made in order for the CCM estimators to be asymptotically conservative. One of these assumptions is that the asymptotic bias caused by confounding variables is positive. This subsection presents a diagnostic test of this assumption in the context of the application presented above.

Figure 3 plots, for the Punishment dependent variable, the $\widehat{ACMET}_1$ and $\widehat{ACMET}_2$ estimates as covariates are added to equation (9) of the estimating regression model.[18] As can

[18]The covariates include respondent age, gender, education, income, employment status, party identification,

be seen, there is a clear downward trend as the number of covariates is increased. This trend illustrates the bias inherent in causal mediation analysis and the gradual reduction of that bias by controlling for confounding variables. Further, it also illustrates in this case that that bias is most likely positive overall since the inclusion of more covariates, which should gradually remove a greater amount of bias, leads to gradual attenuation of the positive estimates.

Figure 3: ACMET Estimates as Additional Covariates are Controlled

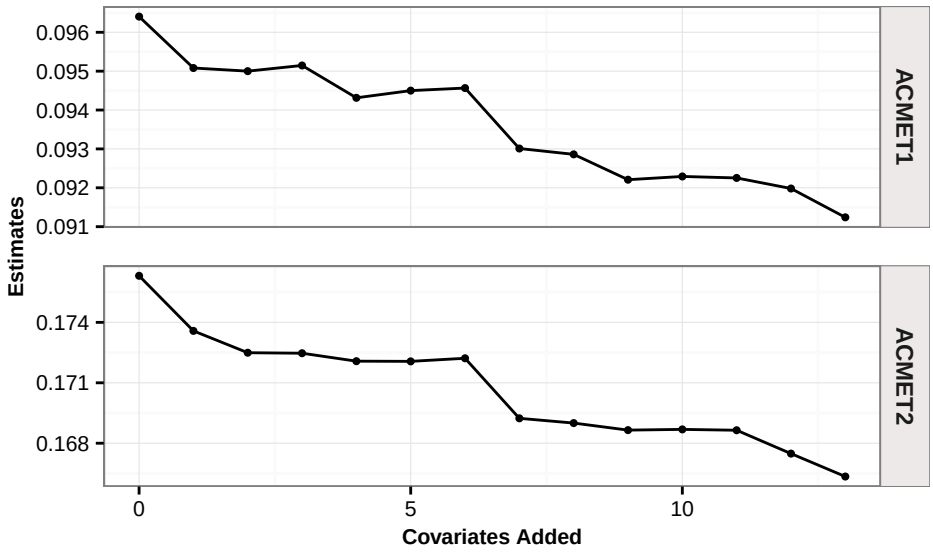

VIII. Conclusions

This paper has introduced a novel set of causal mediation estimands which compare the causal mediation effects of multiple treatments. It has shown that these estimands can be estimated consistently or conservatively under weaker assumptions than can any single average causal mediation effect (ACME). In particular, the usual assumption of no confounding of the mediator-outcome relationship, which is required for consistent estimation of a single ACME, is not necessary in the comparative causal mediation context presented in this paper.

With the gradual accumulation of knowledge and empirical results in various academic subfields and program evaluation contexts, experimental research increasingly involves evaluating multiple treatments—that is, investigating the relative strengths and comparing the causal anatomies of distinct but conceptually or administratively related treatments—rather than simply testing the effects of single treatments. The method of CCM analysis presented in this

voting behavior, military service, and various measures of nationalism.

paper provides a new tool for researchers who are interested in comparing, discovering, and testing the causal mechanism differences between multiple treatments, and would like to do so under the weakest possible set of assumptions.

References

Aaroe, L. (2012). When citizens go against elite directions: Partisan cues and contrast effects on citizens attitudes. *Party Politics*, 18(2):215–233.

Abbott, K. W. and Snidal, D. (2000). Hard and soft law in international governance. *International Organization*, 54(3):421–456.

Angrist, J. D. and Pischke, J.-S. (2009a). *Mostly Harmless Econometrics: An Empiricists Companion*. Princeton: Princeton University Press.

Angrist, J. D. and Pischke, J.-S. (2009b). A note on bias in just identified iv with weak instruments (`http://econ.lse.ac.uk/staff/spischke/mhe/josh/solon_justid_april14.pdf`). Technical report, J. Angrist and J.S. Pischke.

Arceneaux, K. and Kolodny, R. (2009). Educating the least informed: Group endorsements in a grassroots campaign. *American Journal of Political Science*, 53(4):755–770.

Ballard, K. (2007). Convention in peril? riot control agents and the chemical weapons ban. *Arms Control Today*, 37(7).

Baron, R. M. and Kenny, D. A. (1986). The moderator-mediator variable distinction in social psychological research conceptual, strategic, and statistical considerations. *Journal of Personality and Social Psychology*, 51(6):1173–1182.

Bullock, J. G., Green, D. P., and Ha, S. E. (2008). Experimental approaches to mediation: A new guide for assessing causal pathways. Technical report, Yale University.

Bullock, J. G., Green, D. P., and Ha, S. E. (2010). Yes, but whats the mechanism? (dont expect an easy answer). *Journal of Personality and Social Psychology*, 98(4):550–558.

Bullock, J. G. and Ha, S. E. (2011). Mediation analysis is harder than it looks. In Druckman, J. N., Green, D. P., Kuklinski, J. H., and Lupia, A., editors, *Cambridge Handbook of Experimental Political Science*, chapter 35, pages 508–521. Cambridge University Press.

Daniel, R. M., DeStavola, B. L., Cousens, S. N., and Vansteelandt, S. (2015). Causal mediation analysis with multiple mediators. *Biometrics*, 71(1):1–14.

Fieller, E. C. (1954). Some problems in interval estimation. *Journal of the Royal Statistical Society: Series B*, 16(2):175–185.

Franz, V. H. (2007). Ratios: A short guide to confidence limits and proper use (http://arxiv.org/abs/0710.2024). Technical report, arXiv.org.

Gerber, A. S. and Green, D. P. (2012). Mediation. In *Field Experiments: Design, Analysis, and Interpretation*, chapter 10. New York: W. W. Norton & Company.

Gleser, L. J. and Hwang, J. T. (1987). The nonexistence of 100(1-alpha)% confidence sets of finite expected diameter in errors-in-variables and related models. *The Annals of Statistics*, 15(4):1351–1362.

Glynn, A. N. (2012). The product and difference fallacies for indirect effects. *American Journal of Political Science*, 56(1):257–269.

Goren, P., Federico, C. M., and Kittilson, M. C. (2009). Source cues, partisan identities, and political value expression. *American Journal of Political Science*, 53(4):805–820.

Hwang, J. T. G. (1995). Fieller's problems and resampling techniques. *Statistica Sinica*, 5:161–171.

Hyde, S. D. (2015). Experiments in international relations: Lab, survey, and field. *Annual Review of Political Science*, 18:403–424.

Imai, K., Jo, B., and Stuart, E. A. (2011a). Commentary: Using potential outcomes to understand causal mediation analysis. *Multivariate Behavioral Research*, 46(5):842–854.

Imai, K., Keele, L., and Tingley, D. (2010a). A general approach to causal mediation analysis. *Psychological Methods*, 15(4):309–334.

Imai, K., Keele, L., Tingley, D., and Yamamoto, T. (2011b). Unpacking the black box of causality: Learning about causal mechanisms from experimental and observational studies. *American Political Science Review*, 105(4):765–789.

Imai, K., Keele, L., and Yamamoto, T. (2010b). Identification, inference, and sensitivity analysis for causal mediation effects. *Statistical Science*, 25:51–71.

Imai, K. and Yamamoto, T. (2013). Identification and sensitivity analysis for multiple causal mechanisms: Revisiting evidence from framing experiments. *Political Analysis*, 21(2):141–171.

James, L. R. and Brett, J. M. (1984). Mediators, moderators, and tests for mediation. *Journal of Applied Psychology*, 69(2):307–321.

Jo, B. (2008). Causal inference in randomized experiments with mediational processes. *Psychological Methods*, 13:314–336.

Koschat, M. A. (1987). A characterization of the fieller solution. *The Annals of Statistics*, 15(1):462–468.

Loeys, T., Talloen, W., Goubert, L., Moerkerke, B., and Vansteelandt, S. (2016). Assessing moderated mediation in linear models requires fewer confounding assumptions than assessing mediation. *British Journal of Mathematical and Statistical Psychology*, 69(3):352–374.

MacKinnon, D. P., Fairchild, A. J., and Fritz, M. S. (2007). Mediation analysis. *Annual Review of Psychology*, 58:593–614.

Nelson, C. R. and Starz, R. (1988). Some further results on the exact small sample properties of the instrumental variable estimator. Technical report, NBER Technical Working Paper No. 68.

Nicholson, S. P. (2012). Polarizing cues. *American Journal of Political Science*, 56(1):52–66.

Pearl, J. (2001). Direct and indirect effects. Technical report, Proceedings of the 17th Conference on Uncertainty in Artificial Intelligence.

Pearl, J. (2014). Interpretation and identification of causal mediation. *Psychological Methods*, 19(4):459–481.

Preacher, K. J. (2007). Addressing moderated mediation hypotheses: Theory, methods, and prescriptions. *Multivariate Behavioral Research*, 42(1):185–227.

Shpitser, I. and VanderWeele, T. J. (2011). A complete graphical criterion for the adjustment formula in mediation analysis. *The International Journal of Biostatistics*, 7(1).

Slothuus, R. and de Vreese, C. H. (2010). Political parties, motivated reasoning, and issue framing effects. *American Journal of Political Science*, 72(3):630–645.

Tomz, M. (2007). Domestic audience costs in international relations: An experimental approach. *International Organization*, 61:821–840.

VanderWeele, T. J. (2015). *Explanation in Causal Inference: Methods for Mediation and Interaction*. New York: Oxford University Press.

von Luxburg, U. and Franz, V. H. (2009). A geometric approach to confidence sets for ratios: Fieller's theorem, generalizations and boostrap. *Statistica Sinica*, 19:1095–1117.

Supplementary Materials

for

Comparative Causal Mediation: Relaxing the Assumption of No Mediator-Outcome Confounding

Kirk Bansak

Appendix A: Proofs

Proof of Theorem 1.

Call ξ the bias component of $\hat{\beta}$. For notational simplicity, the random estimation variance component of $\hat{\beta}$ will be subsumed into $\dot{\xi}$. Then:

$$\frac{\hat{\alpha}_2\hat{\beta}}{\hat{\alpha}_1\hat{\beta}} = \frac{\hat{\alpha}_2(\beta+\dot{\xi})}{\hat{\alpha}_1(\beta+\dot{\xi})} = \frac{\hat{\alpha}_2\beta+\hat{\alpha}_2\dot{\xi}}{\hat{\alpha}_1\beta+\hat{\alpha}_1\dot{\xi}} = \frac{\hat{\alpha}_2\beta(1+\frac{\dot{\xi}}{\beta})}{\hat{\alpha}_1\beta(1+\frac{\dot{\xi}}{\beta})} = \frac{\hat{\alpha}_2\beta}{\hat{\alpha}_1\beta}$$

Given the exogeneity of T_j (for $j=1,2$), $\hat{\alpha}_j$ is consistent for α_j. Hence:

$$plim\left(\frac{\hat{\alpha}_2\beta}{\hat{\alpha}_1\beta}\right) = \frac{\alpha_2\beta}{\alpha_1\beta} = \frac{ACME_2}{ACME_1}$$

Further:

$$\frac{(\frac{\hat{\alpha}_2\hat{\beta}}{\hat{\tau}_2})}{(\frac{\hat{\alpha}_1\hat{\beta}}{\hat{\tau}_1})} = \frac{(\frac{\hat{\alpha}_2\beta}{\hat{\tau}_2}+\frac{\hat{\alpha}_2\dot{\xi}}{\hat{\tau}_2})}{(\frac{\hat{\alpha}_1\beta}{\hat{\tau}_1}+\frac{\hat{\alpha}_1\dot{\xi}}{\hat{\tau}_1})} = \frac{(\frac{\hat{\alpha}_2\beta}{\hat{\tau}_2}(1+\frac{\dot{\xi}}{\beta}))}{(\frac{\hat{\alpha}_1\beta}{\hat{\tau}_1}(1+\frac{\dot{\xi}}{\beta}))} = \frac{(\frac{\hat{\alpha}_2\beta}{\hat{\tau}_2})}{(\frac{\hat{\alpha}_1\beta}{\hat{\tau}_1})}$$

Given the exogeneity of T_j (for $j=1,2$), $\hat{\alpha}_j$ and $\hat{\tau}_j$ are consistent for α_j and τ_j. Hence:

$$plim\left(\frac{(\frac{\hat{\alpha}_2\beta}{\hat{\tau}_2})}{(\frac{\hat{\alpha}_1\beta}{\hat{\tau}_1})}\right) = \frac{(\frac{\alpha_2\beta}{\tau_2})}{(\frac{\alpha_1\beta}{\tau_1})} = \frac{(\frac{ACME_2}{ATE_2})}{(\frac{ACME_1}{ATE_1})}$$

Proof of Theorem 2.

Consider the following, where the error terms ϵ_j are decomposed into the unobserved confounder variable (or set of variables) X and the true random components ν_j such that $\nu_1 \perp\!\!\!\perp \nu_3$:

$$\begin{aligned} M &= \alpha_1 T_1 + \alpha_2 T_2 + \epsilon_1 \\ &= \alpha_1 T_1 + \alpha_2 T_2 + \psi X + \nu_1 \\ Y &= \delta_1 T_1 + \delta_2 T_2 + \beta M + \gamma_1 T_1 M + \gamma_2 T_2 M + \epsilon_3 \\ &= \delta_1 T_1 + \delta_2 T_2 + \beta M + \gamma_1 T_1 M + \gamma_2 T_2 M + \phi X + \nu_3 \end{aligned}$$

Now consider the following two equations, subsetted by treatment assignment excluding the control condition:

$$\begin{aligned} Y_{T_1=1} &= \delta_1 + (\beta + \gamma_1) M + \epsilon_3^{T_1} \\ &= \delta_1 + \omega_1 M + \epsilon_3^{T_1} \\ Y_{T_2=1} &= \delta_2 + (\beta + \gamma_2) M + \epsilon_3^{T_2} \\ &= \delta_2 + \omega_2 M + \epsilon_3^{T_2} \end{aligned}$$

where the error terms are superscripted to denote that they are the error terms corresponding to particular treatment subsets. Because T_1 and T_2 are both binary, $\hat{\beta} + \hat{\gamma}_j$ estimated from the full interaction equation is identical to $\hat{\omega}_j$ estimated from the subsetted equations, so this subsetting is without loss of generality.

Now, observe that, as with $\hat{\beta}$ in the case of no interaction between T_j and M, $\hat{\omega}_j$ in the subsetted equations are biased estimates of true ω_j. More specifically, similar to before:

$$\begin{aligned} plim\, \hat{\omega}_1 &= \omega_1 + \frac{Cov(\epsilon_1^{T_1}, \epsilon_3^{T_1})}{Var(\epsilon_1^{T_1})} \\ plim\, \hat{\omega}_2 &= \omega_2 + \frac{Cov(\epsilon_1^{T_2}, \epsilon_3^{T_2})}{Var(\epsilon_1^{T_2})} \end{aligned}$$

where

$$\begin{aligned} \epsilon_1^{T_1} &= \psi X^{T_1} + \nu_1^{T_1} \\ \epsilon_1^{T_2} &= \psi X^{T_2} + \nu_1^{T_2} \\ \epsilon_3^{T_1} &= \phi X^{T_1} + \nu_3^{T_1} \\ \epsilon_3^{T_2} &= \phi X^{T_2} + \nu_3^{T_2} \end{aligned}$$

To see this, consider for the subset $T = d$, with the variables demeaned:

$$\begin{aligned} M^{T_d} &= \epsilon_1^{T_d} \\ Y^{T_d} &= \omega_d M^{T_d} + \epsilon_3^{T_d} \end{aligned}$$

leading to:

$$\begin{aligned} plim\, \hat{\omega}_d = (M^{T'_d} M^{T_d})^{-1} M^{T'_d} Y^{T_d} &= (M^{T'_d} M^{T_d})^{-1} M^{T'_d} (\omega_d M^{T_d} + \epsilon_3^{T_d}) \\ &= \omega_d + (M^{T'_d} M^{T_d})^{-1} M^{T'_d} \epsilon_3^{T_d} \\ &= \omega_d + \frac{Cov(M^{T_d}, \epsilon_3^{T_d})}{Var(M^{T_d})} \\ &= \omega_d + \frac{Cov(\epsilon_1^{T_d}, \epsilon_3^{T_d})}{Var(\epsilon_1^{T_d})} \end{aligned}$$

Given the exogeneity of T_j and the status of X as a set of pre-treatment confounders, $f(X|T_1) = f(X|T_2)$, and thus $Cov(\epsilon_1^{T_1}, \epsilon_3^{T_1}) = Cov(\epsilon_1^{T_2}, \epsilon_3^{T_2}) = \psi\phi Var(X)$. However, given the possibility (and virtual certainty with binary M) of heteroskedasticity, $Var(\epsilon_1^{T_1}) \neq Var(\epsilon_1^{T_2})$. Hence:

$$\begin{aligned} plim\, \hat{\omega}_1 &= \omega_1 + \frac{\psi\phi Var(X)}{Var(\epsilon_1^{T_1})} \\ plim\, \hat{\omega}_2 &= \omega_2 + \frac{\psi\phi Var(X)}{Var(\epsilon_1^{T_2})} \end{aligned}$$

Now, call ξ_j the bias component of $\hat{\omega}_j$. Call $\dot{\zeta}_j$ the random estimation variance component of $\hat{\omega}_j$, which of course goes to zero as the sample size grows to infinity. Then:

$$Estimator = \frac{\hat{\alpha}_2\hat{\omega}_2}{\hat{\alpha}_1\hat{\omega}_1} = \frac{\hat{\alpha}_2(\omega_2 + \xi_2 + \dot{\zeta}_2)}{\hat{\alpha}_1(\omega_1 + \xi_1 + \dot{\zeta}_1)} = \frac{\hat{\alpha}_2\omega_2(1 + \frac{\xi_2 + \dot{\zeta}_2}{\omega_2})}{\hat{\alpha}_1\omega_1(1 + \frac{\xi_1 + \dot{\zeta}_1}{\omega_1})} = \frac{\hat{\alpha}_2\omega_2}{\hat{\alpha}_1\omega_1} \cdot \frac{(1 + \frac{\xi_2 + \dot{\zeta}_2}{\omega_2})}{(1 + \frac{\xi_1 + \dot{\zeta}_1}{\omega_1})}$$

$$plim(Estimator) = \frac{\alpha_2\omega_2}{\alpha_1\omega_1} \cdot \frac{(1 + \frac{\xi_2}{\omega_2})}{(1 + \frac{\xi_1}{\omega_1})} = \frac{ACMET_2}{ACMET_1} \cdot Bias$$

Thus, as can be seen, $\frac{\hat{\alpha}_2\hat{\omega}_2}{\hat{\alpha}_1\hat{\omega}_1}$ is not a consistent estimator of the estimand of interest. However, under reasonable and testable conditions, it is consistently conservative (i.e. attenuated toward unity). To see this, consider the asymptotic bias term:

$$Bias = \frac{(1 + \frac{\xi_2}{\omega_2})}{(1 + \frac{\xi_1}{\omega_1})} = \frac{\omega_1\omega_2 + \omega_1\xi_2}{\omega_1\omega_2 + \omega_2\xi_1}$$

Without loss of generality, assume the hypothesis test that $\frac{ACMET_2}{ACMET_1} > 1$ (i.e. $ACMET_2 > ACMET_1$). In such a case, the bias will lead to an underestimate (attenuated toward 1)

of the estimand given two conditions. The first is that the bias term is positive. In the most common case of "consistent mediation"—which is where the total treatment effects and mediation effects are of the same sign, in contrast to "inconsistent mediation" where the mediator actually suppresses the treatment effects (MacKinnon et al., 2007)—this is met when ξ_1 and ξ_2 are also of the same sign as the mediation effects.

The second condition is that $\omega_2\xi_1 > \omega_1\xi_2$, which implies that $\omega_2 Var(\epsilon_1^{T_2}) > \omega_1 Var(\epsilon_1^{T_1})$. This further implies asymptotically (and in expectation) that:

$$\begin{aligned}
\omega_2 Var(\epsilon_1^{T_2}) &> \omega_1 Var(\epsilon_1^{T_1}) \\
\left(\hat{\omega_2} - \frac{\psi\phi Var(X)}{Var(\epsilon_1^{T_2})}\right) Var(\epsilon_1^{T_2}) &> \left(\hat{\omega}_1 - \frac{\psi\phi Var(X)}{Var(\epsilon_1^{T_1})}\right) Var(\epsilon_1^{T_1}) \\
\hat{\omega_2} Var(\epsilon_1^{T_2}) - \psi\phi Var(X) &> \hat{\omega}_1 Var(\epsilon_1^{T_1}) - \psi\phi Var(X) \\
\hat{\omega_2} Var(\epsilon_1^{T_2}) &> \hat{\omega}_1 Var(\epsilon_1^{T_1})
\end{aligned}$$

The same results follow for the estimand $\frac{\left(\frac{ACMET_2}{ATE_2}\right)}{\left(\frac{ACMET_1}{ATE_1}\right)}$.

Appendix B: Proof of Bias Component under No-Interaction Assumption

A full algebraic proof follows below, but for a less formal presentation of this result, consider the following. First, take equations (1) and (3), and partial out T_2, resulting in:

$$M^* = \alpha_1 T_1^* + \epsilon_1^*$$

$$Y^* = \delta_1 T_1^* + \beta M^* + \epsilon_3^*$$

By properties of regression associated with the Frisch-Waugh-Lovell theorem, it can be shown that $\epsilon_i^* = \epsilon_i$ for all i. Hence, we are left with the same sets of relationships and error structures as previously, with $E[\hat{\beta}] = \beta + \frac{cov(\epsilon_1, \epsilon_3)}{var(\epsilon_1)}$. This holds regardless of which treatment we focus on (i.e. regardless of which treatment is partialled out of the regressions).

For a full algebraic proof, consider the following system of equations:

$$M = \alpha_1 X_1 + \alpha_2 X_2 + \epsilon_1 \tag{B10}$$

$$Y = \tau_1 X_1 + \tau_2 X_2 + \epsilon_2 \tag{B11}$$

$$Y = \delta_1 X_1 + \delta_2 X_2 + \beta M + \epsilon_3 \tag{B12}$$

where all variables are demeaned.

For equation B12, call:

$$\tilde{X} = \begin{bmatrix} X_1 & X_2 & M \end{bmatrix} \quad \text{and} \quad \theta = \begin{bmatrix} \delta_1 \\ \delta_2 \\ \beta \end{bmatrix}$$

Then,

$$\hat{\theta} = \begin{bmatrix} \hat{\delta_1} \\ \hat{\delta_2} \\ \hat{\beta} \end{bmatrix} = (\tilde{X}'\tilde{X})^{-1}\tilde{X}'Y = (\tilde{X}'\tilde{X})^{-1}\tilde{X}'(\delta_1 X_1 + \delta_2 X_2 + \beta M + \epsilon_3)$$

$$= \begin{bmatrix} \delta_1 \\ 0 \\ 0 \end{bmatrix} + \begin{bmatrix} 0 \\ \delta_2 \\ 0 \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ \beta \end{bmatrix} + (\tilde{X}'\tilde{X})^{-1}\tilde{X}'\epsilon_3$$

$(\tilde{X}'\tilde{X})^{-1}\tilde{X}'\epsilon_3$ constitutes the bias term, which has three components. The problem is that we cannot simply assume that, in expectation, $(\tilde{X}'\tilde{X})^{-1}\tilde{X}'\epsilon_3 = \mathbf{0}$ because, while X_1 and X_2 are

randomly varied, M is not. The objective pursued below is to characterize the third component, which is the bias on the estimate of β.

First, we know that:

$$\tilde{X}'\tilde{X} = \begin{bmatrix} \sum x_{1i}^2 & \sum x_{1i}x_{2i} & \sum x_{1i}m_i \\ \sum x_{1i}x_{2i} & \sum x_{2i}^2 & \sum x_{2i}m_i \\ \sum x_{1i}m_i & \sum x_{2i}m_i & \sum m_i^2 \end{bmatrix}$$

and

$$\tilde{X}'e_3 = \begin{bmatrix} \sum x_{1i}\epsilon_{3i} \\ \sum x_{2i}\epsilon_{3i} \\ \sum m_i\epsilon_{3i} \end{bmatrix} \approx \begin{bmatrix} 0 \\ 0 \\ \sum m_i\epsilon_{3i} \end{bmatrix}$$

since X_1 and X_2 are randomly varied.

$(\tilde{X}'\tilde{X})^{-1}$ can be found by applying the properties of matrix inversion presented below.

$$(\tilde{X}'\tilde{X})^{-1} = \mathbb{A}^{-1} = \begin{bmatrix} a & b & c \\ d & e & f \\ g & h & i \end{bmatrix}^{-1} = \frac{1}{det(\mathbb{A})} \begin{bmatrix} A & D & G \\ B & E & H \\ C & F & I \end{bmatrix}$$

where

$$\begin{array}{ccc} A = (ei - fh) & D = -(bi - ch) & G = (bf - ce) \\ B = -(di - fg) & E = (ai - cg) & H = -(af - cd) \\ C = (dh - eg) & F = -(ah - bg) & I = (ae - bd) \end{array}$$

and

$$det(\mathbb{A}) = aA + bB + cC$$

The result is that the bias term is:

$$(\tilde{X}'\tilde{X})^{-1}\tilde{X}'\epsilon_3 = \frac{1}{det(\tilde{X}'\tilde{X})} \begin{bmatrix} G\sum m_i\epsilon_{3i} \\ H\sum m_i\epsilon_{3i} \\ I\sum m_i\epsilon_{3i} \end{bmatrix}$$

$$= \frac{1}{aA + bB + cC} \begin{bmatrix} (bf - ce)\sum m_i\epsilon_{3i} \\ -(af - cd)\sum m_i\epsilon_{3i} \\ (ae - bd)\sum m_i\epsilon_{3i} \end{bmatrix}$$

which means that the additive bias on $\hat{\beta}$ is:

$$\frac{1}{aA + bB + cC}(ae - bd)\sum m_i\epsilon_{3i}$$

To characterize this bias, we need characterize each component and furthermore to use equation B10 to replace M.

$$\sum m_i\epsilon_{3i} = \sum(\alpha_1 x_{1i} + \alpha_2 x_{2i} + \epsilon_{1i})\epsilon_{3i} = 0 + 0 + \sum \epsilon_{1i}\epsilon_{3i} = cov(\epsilon_1, \epsilon_3)$$

since $E[\epsilon_1] = 0$ and $E[\epsilon_3] = 0$ by assumption.

$$ae - bd = \sum x_{1i}^2 \sum x_{2i}^2 - \sum x_{1i}x_{2i} \sum x_{1i}x_{2i} = var(X_1)var(X_2) - cov(X_1, X_2)^2$$

since X_1 and X_2 are both demeaned.

$$aA + bB + cC$$

$$\begin{aligned}
&= \sum x_{1i}^2 \sum x_{2i}^2 \sum m_i^2 - \sum x_{1i}^2 \sum x_{2i}m_i \sum x_{2i}m_i \\
&+ \sum x_{1i}x_{2i} \sum x_{2i}m_i \sum x_{1i}m_i - \sum x_{1i}x_{2i} \sum x_{1i}x_{2i} \sum m_i^2 \\
&+ \sum x_{1i}m_i \sum x_{1i}x_{2i} \sum x_{2i}m_i - \sum x_{1i}m_i \sum x_{2i}^2 \sum x_{1i}m_i
\end{aligned}$$

and substituting in $m_i = \alpha_1 x_{1i} + \alpha_2 x_{2i} + \epsilon_{1i}$, with a bit of algebraic simplification, it can be shown that

$$\begin{aligned}
aA + bB + cC &= \left[\sum \epsilon_{1i}^2\right] \left[\sum x_{1i}^2 \sum x_{2i}^2 - \sum x_{1i}x_{2i} \sum x_{1i}x_{2i}\right] \\
&= var(\epsilon_1) \left(var(X_1)var(X_2) - cov(X_1, X_2)^2\right)
\end{aligned}$$

Hence,

$$\begin{aligned}
&\frac{1}{aA + bB + cC}(ae - bd) \sum m_i \epsilon_{3i} \\
&= \frac{1}{var(\epsilon_1) \left(var(X_1)var(X_2) - cov(X_1, X_2)^2\right)} \left(var(X_1)var(X_2) - cov(X_1, X_2)^2\right) cov(\epsilon_1, \epsilon_3) \\
&= \frac{cov(\epsilon_1, \epsilon_3)}{var(\epsilon_1)}
\end{aligned}$$

Appendix C: Finite-Sample Adjustments

Finite-sample adjustments for the CCM estimators can be derived using Taylor series expansion.

Consider the first estimator under the no-interaction assumption, $\frac{\hat{\alpha}_2\hat{\beta}}{\hat{\alpha}_1\hat{\beta}}$, which can (quite apparently) be simplified to $\frac{\hat{\alpha}_2}{\hat{\alpha}_1}$. Similarly, the estimand of interest can be seen simply as:

$$\frac{ACME_2}{ACME_1} = \frac{\alpha_2\beta}{\alpha_1\beta} = \frac{\alpha_2}{\alpha_1} = \frac{E[\hat{\alpha}_2]}{E[\hat{\alpha}_1]}$$

However, a first problem is the following:

$$E\left[\frac{\hat{\alpha}_2}{\hat{\alpha}_1}\right] \neq \frac{E[\hat{\alpha}_2]}{E[\hat{\alpha}_1]} = \frac{\alpha_2}{\alpha_1}$$

A second problem is that $E\left[\frac{\hat{\alpha}_2}{\hat{\alpha}_1}\right]$ may not even exist. To address both of these problems, the estimator $\frac{\hat{\alpha}_2}{\hat{\alpha}_1}$, which will be denoted as $f(\hat{\Theta})$ can be approximated using a (second-order) multivariate Taylor series expansion around the estimand $f(\Theta)$:

$$f(\hat{\Theta}) \approx f(\Theta) + \sum_{\theta\in\Theta}(\hat{\theta}-\theta)f_{\hat{\theta}}(\Theta) + \frac{1}{2}\sum_{\theta\in\Theta}\sum_{\theta'\in\Theta}(\hat{\theta}-\theta)(\hat{\theta}'-\theta')f_{\hat{\theta}\hat{\theta}'}(\Theta)$$

where Θ contains the full set of parameters (denoted individually by θ), $f_{\hat{\theta}}$ refers to the first derivative of f with respect to $\hat{\theta}$, and $f_{\hat{\theta}\hat{\theta}'}$ refers to the second derivative of f with respect to $\hat{\theta}$ and $\hat{\theta}'$.

If we treat the higher-order terms in the Taylor series expansion as negligible, as conventionally done, then we can identify the approximate divergence between the estimator and the estimand, which is a quantity for which we can characterize the moments:

$$E\left[\sum_{\theta\in\Theta}(\hat{\theta}-\theta)f_{\hat{\theta}}(\Theta) + \frac{1}{2}\sum_{\theta\in\Theta}\sum_{\theta'\in\Theta}(\hat{\theta}-\theta)(\hat{\theta}'-\theta')f_{\hat{\theta}\hat{\theta}'}(\Theta)\right]$$

The first-order terms in this expression are zero in expectation (i.e. $E[\hat{\theta}-\theta]=0$), while the leading components of the second-order terms are covariances in expectation (i.e. $E[(\hat{\theta}-\theta)(\hat{\theta}'-\theta')] = Cov(\hat{\theta},\hat{\theta}')$). Thus, the divergence is approximately:

$$\frac{1}{2}\sum_{\theta\in\Theta}\sum_{\theta'\in\Theta}Cov(\hat{\theta},\hat{\theta}')f_{\hat{\theta}\hat{\theta}'}(\Theta)$$

This divergence can thus be estimated—by plugging in $\hat{\Theta}$ for Θ and estimating the covariances—and then subtracted from the simple estimator $f(\hat{\Theta})$ of interest to yield an adjusted estimator that is approximately centered on the estimand of interest. Also evident from the expression is that this divergence term goes to zero as the sample size n grows to infinity. The following applies this process to the actual estimators in question.

A. Adjusted Estimators under the No-Interaction Assumption

A.1. Adjusted Estimator 1

The estimator for the first estimand is $\frac{\hat{\alpha}_2\hat{\beta}}{\hat{\alpha}_1\hat{\beta}} = \frac{\hat{\alpha}_2}{\hat{\alpha}_1}$. In expectation, the second-order Taylor Series expansion of the estimator, $T(\frac{\hat{\alpha}_2}{\hat{\alpha}_1})$, around the estimand is:

$$E\left[T\left(\frac{\hat{\alpha}_2}{\hat{\alpha}_1}\right)\right] \approx \frac{\alpha_2}{\alpha_1} - \frac{Cov(\hat{\alpha}_1, \hat{\alpha}_2)}{\alpha_1^2} + \frac{Var(\hat{\alpha}_1)\alpha_2}{\alpha_1^3}$$

Hence, we can identify the component of the approximation that diverges from the estimand. Because of the exogeneity of T, α_1 and α_2 can both be estimated without bias, allowing for the individual pieces of that component to be estimated by regression. This can then be subtracted from the estimator $\frac{\hat{\alpha}_2}{\hat{\alpha}_1}$ to yield an adjusted estimator approximately centered on the estimand:

$$\frac{\hat{\alpha}_2}{\hat{\alpha}_1} + \frac{\widehat{Cov}(\hat{\alpha}_1, \hat{\alpha}_2)}{\hat{\alpha}_1^2} - \frac{\widehat{Var}(\hat{\alpha}_1)\hat{\alpha}_2}{\hat{\alpha}_1^3}$$

In the special case of balanced control and treatment assignment (i.e. $P(C) = P(T_1) = P(T_2) = \frac{1}{3}$), the adjusted estimator simplifies to:

$$\frac{\hat{\alpha}_2}{\hat{\alpha}_1} + \frac{3\hat{\sigma}_{e_1}^2}{\hat{\alpha}_1^2 n} - \frac{6\hat{\sigma}_{e_1}^2\hat{\alpha}_2}{\hat{\alpha}_1^3 n}$$

where $\hat{\sigma}_{e_1}^2$ refers to the estimated error variance from equation (1). Clearly, as n grows to infinity, this converges on the simple estimator $\frac{\hat{\alpha}_2}{\hat{\alpha}_1}$.

A.2. Adjusted Estimator 2

The simple estimator for the second estimand is $\frac{(\frac{\hat{\alpha}_2\hat{\beta}}{\hat{\tau}_2})}{(\frac{\hat{\alpha}_1\hat{\beta}}{\hat{\tau}_1})} = \frac{(\frac{\hat{\alpha}_2}{\hat{\tau}_2})}{(\frac{\hat{\alpha}_1}{\hat{\tau}_1})} = \frac{\hat{\alpha}_2\hat{\tau}_1}{\hat{\alpha}_1\hat{\tau}_2}$. As above, a second-order Taylor Series expansion can be used to formulate an adjusted estimator that is approximately centered on the estimand in finite samples:

$$\begin{aligned}
&\frac{\hat{\alpha}_2\hat{\tau}_1}{\hat{\alpha}_1\hat{\tau}_2} - \widehat{Var}(\hat{\alpha}_1)\frac{\hat{\alpha}_2\hat{\tau}_1}{\hat{\alpha}_1^3\hat{\tau}_2} - \widehat{Var}(\hat{\tau}_2)\frac{\hat{\alpha}_2\hat{\tau}_1}{\hat{\alpha}_1\hat{\tau}_2^3} + \widehat{Cov}(\hat{\alpha}_2, \hat{\alpha}_1)\frac{\hat{\tau}_1}{\hat{\alpha}_1^2\hat{\tau}_2} + \widehat{Cov}(\hat{\alpha}_2, \hat{\tau}_2)\frac{\hat{\tau}_1}{\hat{\alpha}_1\hat{\tau}_2^2} \\
&-\widehat{Cov}(\hat{\alpha}_2, \hat{\tau}_1)\frac{1}{\hat{\alpha}_1\hat{\tau}_2} - \widehat{Cov}(\hat{\alpha}_1, \hat{\tau}_2)\frac{\hat{\alpha}_2\hat{\tau}_1}{\hat{\alpha}_1^2\hat{\tau}_2^2} + \widehat{Cov}(\hat{\alpha}_1, \hat{\tau}_1)\frac{\hat{\alpha}_2}{\hat{\alpha}_1^2\hat{\tau}_2} + \widehat{Cov}(\hat{\tau}_2, \hat{\tau}_1)\frac{\hat{\alpha}_2}{\hat{\alpha}_1\hat{\tau}_2^2}
\end{aligned}$$

B. Adjusted Estimators when Relaxing the No-Interaction Assumption

Having discarded the no-interaction assumption, the estimator of the first estimand of interest ($\frac{ACMET_2}{ACMET_1}$) is $\frac{\hat{\alpha}_2(\hat{\beta}+\hat{\gamma}_2)}{\hat{\alpha}_1(\hat{\beta}+\hat{\gamma}_1)} = \frac{\hat{\alpha}_2\hat{\omega}_2}{\hat{\alpha}_1\hat{\omega}_1}$.

As shown, $plim\ \hat{\omega}_j = \omega_j + \xi_j$, because of a confounding bias that does not disappear asymptotically. Let ω_j^* denote the biased and inconsistent version of ω_j (i.e. $plim\ \hat{\omega}_j = \omega_j^*$). As shown above, under certain reasonable and testable assumptions, $\frac{\alpha_2\omega_2^*}{\alpha_1\omega_1^*}$ is conservative (i.e. attenuated toward 1) for $\frac{\alpha_2\omega_2}{\alpha_1\omega_1}$ and hence the estimator of interest is asymptotically conservative for the estimand of interest. Unfortunately, for two reasons, this does not mean that in small samples the estimator of interest is in expectation also conservative. First, as before, the expectation may not even actually exist. Second, also as before, the ratio form of the estimand leads the estimator to be decentered from the point to which it converges. However, also as in the case with the no-interaction assumption, a second-order Taylor Series expansion can be used to construct an adjusted estimator that in finite samples is approximately centered upon the conservative point for which the estimator is consistent.

Specifically, the adjusted estimator is:

$$\begin{aligned}
&\frac{\hat{\alpha}_2\hat{\omega}_2}{\hat{\alpha}_1\hat{\omega}_1} - \widehat{Var}(\hat{\alpha}_1)\frac{\hat{\alpha}_2\hat{\omega}_2}{\hat{\alpha}_1^3\hat{\omega}_1} - \widehat{Var}(\hat{\omega}_1)\frac{\hat{\alpha}_2\hat{\omega}_2}{\hat{\alpha}_1\hat{\omega}_1^3} + \widehat{Cov}(\hat{\alpha}_2,\hat{\alpha}_1)\frac{\hat{\omega}_2}{\hat{\alpha}_1^2\hat{\omega}_1} + \widehat{Cov}(\hat{\alpha}_2,\hat{\omega}_1)\frac{\hat{\omega}_2}{\hat{\alpha}_1\hat{\omega}_1^2} \\
&-\widehat{Cov}(\hat{\alpha}_2,\hat{\omega}_2)\frac{1}{\hat{\alpha}_1\hat{\omega}_1} - \widehat{Cov}(\hat{\alpha}_1,\hat{\omega}_1)\frac{\hat{\alpha}_2\hat{\omega}_2}{\hat{\alpha}_1^2\hat{\omega}_1^2} + \widehat{Cov}(\hat{\alpha}_1,\hat{\omega}_2)\frac{\hat{\alpha}_2}{\hat{\alpha}_1^2\hat{\omega}_1} + \widehat{Cov}(\hat{\omega}_1,\hat{\omega}_2)\frac{\hat{\alpha}_2}{\hat{\alpha}_1\hat{\omega}_1^2}
\end{aligned}$$

where $\hat{\omega}_j = \hat{\beta} + \hat{\gamma}_j$ from Equation 9, and the covariance terms can be estimated via the bootstrap.

Following the same approach for the second CCM estimand, the adjusted version of the second estimator, $\frac{\left(\frac{\hat{\alpha}_2\hat{\omega}_2}{\hat{\tau}_2}\right)}{\left(\frac{\hat{\alpha}_1\hat{\omega}_1}{\hat{\tau}_1}\right)}$, is:

$$\begin{aligned}
&\frac{\left(\frac{\hat{\alpha}_2\hat{\omega}_2}{\hat{\tau}_2}\right)}{\left(\frac{\hat{\alpha}_1\hat{\omega}_1}{\hat{\tau}_1}\right)} - \widehat{Var}(\hat{\alpha}_1)\frac{\hat{\alpha}_2\hat{\omega}_2\hat{\tau}_1}{\hat{\alpha}_1^3\hat{\omega}_1\hat{\tau}_2} - \widehat{Var}(\hat{\omega}_1)\frac{\hat{\alpha}_2\hat{\omega}_2\hat{\tau}_1}{\hat{\alpha}_1\hat{\omega}_1^3\hat{\tau}_2} - \widehat{Var}(\hat{\tau}_2)\frac{\hat{\alpha}_2\hat{\omega}_2\hat{\tau}_1}{\hat{\alpha}_1\hat{\omega}_1\hat{\tau}_2^3} \\
&+\widehat{Cov}(\hat{\alpha}_2,\hat{\alpha}_1)\frac{\hat{\omega}_2\hat{\tau}_1}{\hat{\alpha}_1^2\hat{\omega}_1\hat{\tau}_2} - \widehat{Cov}(\hat{\alpha}_2,\hat{\omega}_2)\frac{\hat{\tau}_1}{\hat{\alpha}_1\hat{\omega}_1\hat{\tau}_2} + \widehat{Cov}(\hat{\alpha}_2,\hat{\omega}_1)\frac{\hat{\omega}_2\hat{\tau}_1}{\hat{\alpha}_1\hat{\omega}_1^2\hat{\tau}_2} \\
&+\widehat{Cov}(\hat{\alpha}_2,\hat{\tau}_2)\frac{\hat{\omega}_2\hat{\tau}_1}{\hat{\alpha}_1\hat{\omega}_1\hat{\tau}_2^2} - \widehat{Cov}(\hat{\alpha}_2,\hat{\tau}_1)\frac{\hat{\omega}_2}{\hat{\alpha}_1\hat{\omega}_1\hat{\tau}_2} + \widehat{Cov}(\hat{\alpha}_1,\hat{\omega}_2)\frac{\hat{\alpha}_2\hat{\tau}_1}{\hat{\alpha}_1^2\hat{\omega}_1\hat{\tau}_2} \\
&-\widehat{Cov}(\hat{\alpha}_1,\hat{\omega}_1)\frac{\hat{\alpha}_2\hat{\omega}_2\hat{\tau}_1}{\hat{\alpha}_1^2\hat{\omega}_1^2\hat{\tau}_2} - \widehat{Cov}(\hat{\alpha}_1,\hat{\tau}_2)\frac{\hat{\alpha}_2\hat{\omega}_2\hat{\tau}_1}{\hat{\alpha}_1^2\hat{\omega}_1\hat{\tau}_2^2} + \widehat{Cov}(\hat{\alpha}_1,\hat{\tau}_1)\frac{\hat{\alpha}_2\hat{\omega}_2}{\hat{\alpha}_1^2\hat{\omega}_1\hat{\tau}_2} \\
&+\widehat{Cov}(\hat{\omega}_2,\hat{\omega}_1)\frac{\hat{\alpha}_2\hat{\tau}_1}{\hat{\alpha}_1\hat{\omega}_1^2\hat{\tau}_2} + \widehat{Cov}(\hat{\omega}_2,\hat{\tau}_2)\frac{\hat{\alpha}_2\hat{\tau}_1}{\hat{\alpha}_1\hat{\omega}_1\hat{\tau}_2^2} - \widehat{Cov}(\hat{\omega}_2,\hat{\tau}_1)\frac{\hat{\alpha}_2}{\hat{\alpha}_1\hat{\omega}_1\hat{\tau}_2} \\
&-\widehat{Cov}(\hat{\omega}_1,\hat{\tau}_2)\frac{\hat{\alpha}_2\hat{\omega}_2\hat{\tau}_1}{\hat{\alpha}_1\hat{\omega}_1^2\hat{\tau}_2^2} + \widehat{Cov}(\hat{\omega}_1,\hat{\tau}_1)\frac{\hat{\alpha}_2\hat{\omega}_2}{\hat{\alpha}_1\hat{\omega}_1^2\hat{\tau}_2} + \widehat{Cov}(\hat{\tau}_2,\hat{\tau}_1)\frac{\hat{\alpha}_2\hat{\omega}_2}{\hat{\alpha}_1\hat{\omega}_1\hat{\tau}_2^2}
\end{aligned}$$

In sum, if the assumption of no interaction between the treatments and the mediator is relaxed, the CCM estimators are no longer consistent, but they are asymptotically conservative

provided additional conditions are met. Those additional conditions are both theoretically reasonable and empirically testable. Furthermore, finite-sample adjustments can be added to the estimators such that they are also conservative in smaller samples.

Appendix D: Simulations when Relaxing the No-Interaction Assumption

To illustrate the properties of the CCM estimators once the no-interaction assumption has been relaxed, this section presents the results of a simulation. The data-generating process was similar to that of the simulation presented earlier except, in this case, the effect of the mediator on the outcome involves interactions with both treatments. In addition, the simulated sample size has been increased to 1000 units per treatment condition in order to better illustrate the asymptotic tendencies.[19] As before, positive bias is introduced by construction through the omission in the estimation of a confounder that affects both the outcome and mediator. Also as before, the ACME for the treated for the second treatment is larger than that of the first treatment; further, the interaction between the mediator and the second treatment is also made larger than the interaction between the mediator and the first treatment. Thus, the additional conditions required for conservative estimation of the CCM estimands are met. Figure D1 shows the resulting estimates in the simulation.

Figure D1: Comparative Causal Mediation Simulation: With Interactions

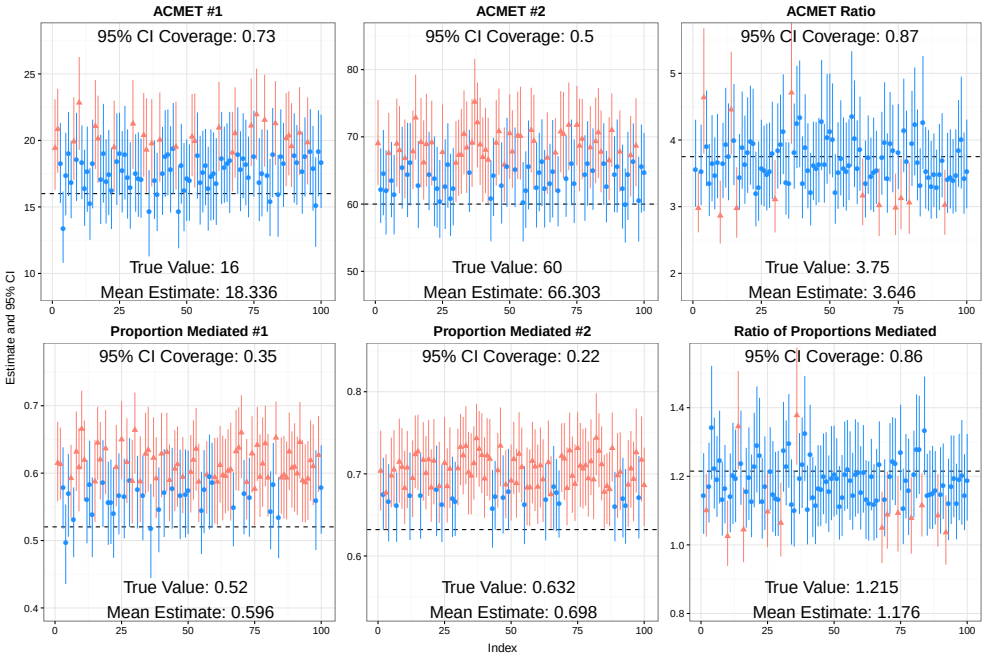

As can be seen in the top row of Figure D1, the estimators of the ACMEs for the treated are again biased upward and, as a result, also have bad confidence-interval coverage. In

[19] For this reason, the finite-sample adjustments make little difference, and hence the adjusted estimators are not presented here.

contrast, however, the estimator of the ratio of ACMEs for the treated is much more well-behaved. While no longer consistent, and hence not properly centered in this medium-sized sample, the estimator is conservative (attenuated toward unity), as indicated by the mean estimate being closer to one than the true value. As a result of this conservatism, there is unfortunately confidence-interval under-coverage. However, what makes this problem less concerning is that the under-coverage is the result of attenuated estimates, as shown by the majority of bad confidence intervals being below the true value, rather than the result of systematically undersized confidence intervals.

The results are similar for the bottom row of Figure D1, which presents the estimates for the proportions mediated, as well as the ratio of the proportions mediated. Again, the traditional estimators are biased upward, while the CCM estimator is conservative.

Appendix E: Application Text

Prologue

Please consider the following hypothetical scenario:
ISIS militants in Iraq were threatening rocket attacks on neighboring countries in the region. In response, the U.S. government considered taking military action. The U.S. ruled out drone strikes and other options because the ISIS militants were hiding in a civilian zone, and the U.S. government wanted to avoid harming civilians. Instead, U.S. commandos were deployed in a covert operation. In order to avoid inflicting permanent harm on nearby civilians, the commandos used a non-lethal "incapacitating" chemical gas to knock out and capture the ISIS militants. However, critics of the operation have pointed out that people have varying levels of sensitivity to the incapacitating gas, and exposure can be fatal for some people. Hence, the operation may have put civilian lives in harm's way.

Treatment

CONTROL (*no additional information provided*)

OR

NON-LEGAL TREATMENT: Furthermore, the U.S. government has pledged never to use incapacitating chemical gas in previous public statements. Hence, the U.S. government has broken its pledge.

OR

LEGAL TREATMENT: Furthermore, the U.S. government has pledged never to use incapacitating chemical gas under its membership in the Chemical Weapons Convention, the international treaty banning chemical weapons. Hence, the U.S. government has broken international law.

DV 1: Disapproval

In general, do you approve or disapprove of the U.S. government's decision to use the incapacitating gas in the operation?

- *Approve Strongly, Approve, Neither Approve nor Disapprove, Disapprove, Disapprove Strongly*
- Variable is dichotomized for analysis, with 1 indicating "Disapprove" or "Disapprove Strongly," and 0 otherwise.

DV 2: Punishment

Imagine that one of your U.S. Senators voted in favor of using the incapacitating gas. Would this increase or decrease your willingness to vote for that Senator in the next election?

- *Increase Greatly, Increase, Neither Increase nor Decrease, Decrease, Decrease Greatly*
- Variable is dichotomized for analysis, with 1 indicating "Decrease" or "Decrease Greatly," and 0 otherwise.

Mediator: Perceived Immorality

To what extent do you believe that the decision to use the incapacitating gas in the operation was morally right or wrong?

- *Definitely Right, Probably Right, Not Morally Right or Wrong, Probably Wrong, Definitely Wrong*
- Variable is dichotomized for analysis, with 1 indicating "Probably Wrong" or "Definitely Wrong" and 0 otherwise.

Appendix F: Application Sample Demographics

Table F1: Sample Demographics

Gender	
Female	Male
46.3%	53.7%

Age			
18-29	30-44	45-64	65+
38.9%	41.1%	18.5%	1.6%

Education			
No High School	High School	Some College	College Graduate
0.8%	11.9%	34.1%	53.1%